Estimating Swarm Parameters by Evolutionary Learning

Richard D. Coates
Department of Computer Science
University of York, York YO10 5GH, UK

Simon J. Hickinbotham
Department of Computer Science
University of York, York YO10 5GH, UK
email: sjh@cs.york.ac.uk

ABSTRACT

If we assume that the collective dynamics of wild animals can be modelled, it would be desirable to recover the dynamics of the model via interaction with them. In this paper, we demonstrate that it is possible to recover the parameters of a shoaling model used by a swarm. This can be achieved by evolving the parameters of a single agent that interacts with the swarm. We present an evaluation of this approach, using a genetic algorithm to learn the parameters of a shoaling model.

Categories and Subject Descriptors

D2.2 [**SOFTWARE ENGINEERING**]: Design Tools and Techniques—*Evolutionary prototyping*

General Terms

Experimentation

Keywords

Swarm robotics, Genetic algorithms, Fitness evaluation, Parameter tuning

1. INTRODUCTION

In nature many animals travel in swarms. Several models have been proposed to simulate this phenomena [1]. In each of these models, individuals follow local rules which produce the swarm as an emergent phenomenon. In this paper we explore the idea that can recover the parameters of a swarm by monitoring its interactions with the rest of the swarm. This approach was inspired by the work of Faria et al [2], in which a robotic fish (or *robofish*) interacted with a shoal of sticklebacks *Gasterosteus aculeatus* (L) (or *modelfish*).

It is only possible to determine whether model parameters can be learned in this way if the behaviour of the swarm is already known to follow a fully parameterised model. Accordingly, we test the approach in simulation, where a swarm of modelfish follow a pre-specified model, but the parameters are unknown to a simulated robofish.

The experiments we report here used a two-dimensional version of the Couzin [1] model of shoaling behaviour as the basis of simulated fish movement based on zones of perception around each fish. The presence of neighbours within these zones governs the fish's future movement. We attempt to recover the widths of the zones of Attraction (ZOA), Orientation (ZOO) and Repulsion (ZOR) of the modelfish via interaction with a series robofish with an evolved set of zone widths.

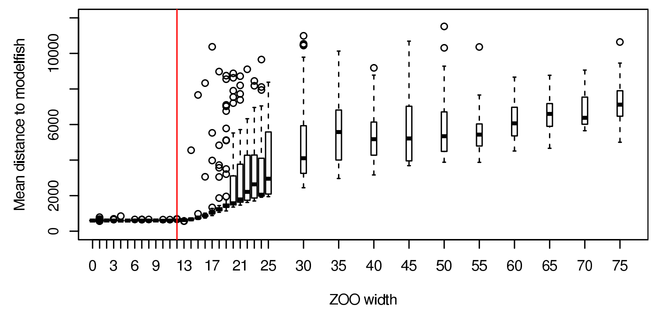

Figure 1: Effect on fitness of changes in width of ZOO. The vertical red line indicates the target width value used by the modelfish.

2. METHODOLOGY

Our experiments used the average Euclidean distance from the robofish to each of the modelfish as the fitness function for a genetic algorithm (GA). This fitness function is based on two assumptions: (1), that similar, but not identical behaviour will allow the robofish to interact with the swarm; (2) that only one configuration of the swarm model will induce this behaviour. A single robofish was used in each simulation to ensure that the robofish's behaviour was determined only by the model it was running and the modelfish. The parameters used by the model fish in the swarm was set throughout the experiments with parameters used in [3].

For all the experiments the number of fish in each trial (inclusive of the robofish) was set to 100. The swarm was run for 5000 time steps before 5000 time steps of monitored behaviour. Here the state of the system was sampled every 50 time steps. In addition to this each configuration was performed 5 times to further reduce the variability in runs.

ZOA: It would be expected that a high ZOA would allow the robofish to find the swarm more easily if it became separated from the group. We found that even a small ZOA with a width of 1 is effective in maintaining the robofish's contact with the swarm. This suggests that the ZOA has a role in limiting the chance of *escape* from the swarm, rather

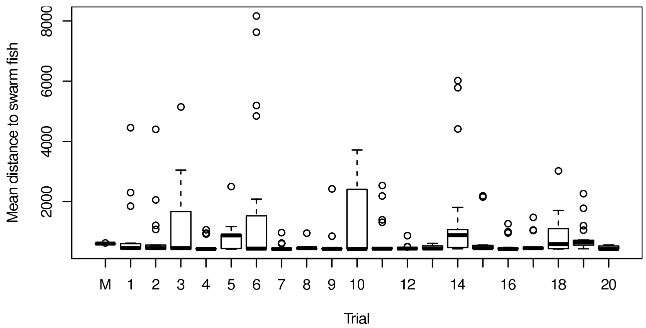

Figure 2: Distribution of fitness at end of 50 generations of 20 trials of the GA. Column 'M' indicates the distribution of fitness for the modelfish

than directing behaviour to *seek* the swarm when swimming alone. Where a ZOA is present in a trial, the swarm has usually formed and contains the robofish by the time fitness measurements commence. Note also that our fitness measure only samples the position of the robofish after a period of 5000 time steps. If we began to measure fitness from the initialisation period, we may see the effects of varying the ZOA during the process of swarm formation.

ZOO: The ZOO has a much clearer effect on the fitness of the robofish, as shown in figure 1. At low values for ZOO the robofish performs very well. It is possible that the interplay between ZOA and ZOO (described above) is the reason for this. If ZOO is disabled the ZOA will still drive the fish towards the swarm. The mean distance between the robofish and the model fish increases when the robofish ZOO width rises above the ZOO width of the model fish. Between 14 and 19 a slow phase change occurs from a fairly uniform fitness to a more variable, unfit behaviour afterwards. We suggest that this phase change is a by-product of the behaviour a larger ZOO induces: A situation arises in which the model fish are in the ZOO of the robofish, but the robofish is in the ZOA of a few modelfish. Since other modelfish are also in the ZOO of the modelfish, influence of the robofish on the shoal is negligible. As the width of the ZOO increases, their is an increasing likelihood that the robofish will leave the swarm.

3. EVOLVING THE ROBOFISH

The previous section demonstrated a relationship between changes in the individual model widths and mean distance to the shoal fish. In this section, we show how a GA can be used to find combinations of widths for the ZOA, ZOO and ZOR that minimises the mean distance to the swarm. We maintained a population of 20 robofish models per generation throughout the trial. Crossover was set to 90%, and mutation (occurring with a 5% chance) changed the Zone widths following a Gaussian distribution with variance of 5% of the current value. The three zone widths were initialised with random integers in the range [0,75]. The GA was run for 50 generations using tournament selection. This was repeated 20 times.

Figure 2 shows the final distribution of mean distance to swarm fish for the robofish population in the final generation of the 20 runs of the GA. The mean distance to shoal fish for robofish with the same width values as the model fish is shown in the column marked 'M'. It can be seen that trials 3, 5, 6, 10, 14 and 18 have not fully converged, but the other fourteen trials show that the GA has successfully reduced the mean distance to the swarm fish, as specified by the fitness function. Those trials which did not converge have a mixture of individuals with low and high mean distances to the modelfish. It is likely that these runs would eventually converge to low average mean distances across the population if allowed to run for longer. Note that the mean distance to shoal fish was higher in the control robofish 'M' that used the model zone widths. This reveals an issue with the fitness function - it was designed to evolve a fish that interacted with the shoal, but there is nothing in the fitness function to induce the evolved robofish to *mimic* the behaviour of the model fish. This is why the mean distance is reduced to a minimum, rather than converging on the value that the modelfish parameters generate.

4. CONCLUSIONS

We have investigated ways of recovering the underlying model parameters of the swarm indirectly, via the interactions of an individual with the swarm. Our initial trials revealed the following observations: a ZOA is needed to produce a swarm (i.e. the ZOA width must be greater than zero), but the size of the ZOA makes no difference once the swarm is formed; and the larger the ZOO the wider the swarm distribution, since individuals can influence the direction of the swarm whilst remaining relatively widely dispersed.

The fitness measure was based on the distances between the robofish and the modelfish and could be improved by reference to the distance of modelfish *to each other.* If this technique were to be used with real fish, it would be possible to do this using computer vision techniques, although issues with sampling time might arise.

5. ACKNOWLEDGMENTS

The authors thank A. Jamie Wood for provision of source code for the swarm model. Simon Hickinbotham is funded by the Plazzmid project, EPSRC grant EP/F031033/1.

6. REFERENCES

[1] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *J. Theor. Biol.*, 218:1–11, 2002.

[2] J.J. Faria, J.R.G. Dyer, R.O. ClÃľment, I.D. Couzin, N Holt, A.J.W. Ward, D. Waters, and J. Krause. A novel method for investigating the collective behaviour of fish: introducing 'robofish'. *Behav. Ecol. Sociobiol.*, 64(8):1211–1218, 2010.

[3] A. J. J. Wood and G. J. J. Ackland. Evolving the selfish herd: emergence of distinct aggregating strategies in an individual-based model. *Proc. Roy. Soc. B*, 274:1637–1642, May 2007.