

(Not) Evolving Collective Behaviours in Synthetic Fish

Nahum Zaera¹, Dave Cliff², and Janet Bruten¹

¹Hewlett-Packard Laboratories Bristol,
Filton Road, Stoke Gifford, Bristol BS12 6QZ, U.K.

²School of Cognitive and Computing Sciences,
University of Sussex, Brighton BN1 9QH, U.K.

nz@hplb.hpl.hp.com davec@cogs.susx.ac.uk jlb@hplb.hpl.hp.com

Abstract

We describe a series of experiments in which artificial evolution is used to develop neural-network sensory-motor controllers for software animats which exhibit collective movement behaviours in three dimensions. We successfully evolved controllers which displayed simple group behaviours such as dispersal and aggregation. However, attempts to evolve realistic-looking schooling behaviours never succeeded. The problem appears to be due to the difficulty of formulating an evaluation function which captures what schooling is. We argue that formulating an effective fitness evaluation function for use in evolving controllers can be at least as difficult as hand-crafting an effective controller design. Although our paper concentrates on schooling, we believe that this is likely to be a general issue, and is a serious problem which can be expected to be experienced over a variety of problem domains.

1 Introduction

Artificial evolution, in the form of genetic algorithms (GAs), is often advocated as a labour-saving approach, where the design of complex artefacts can be achieved semi-automatically. Several authors have reported work in which an evolutionary process has been used to develop designs for sensory-motor “controller” coordination mechanisms for animats (i.e. artificial autonomous agents, either mobile robots or virtual “software creatures”).

In a recent review of selected work in this area [MAT95], it was noted that in many cases the behaviours produced by the final evolved controllers were relatively simple, and that controllers which produce equivalent behaviours could feasibly be designed manually with no more effort than was required to construct the artificial evolution process. It is argued that such results are

to be expected at this early stage in the development of the field, where the concept of using artificial evolution is being demonstrated, and where the evolutionary techniques are being developed, refined, and extended.

In this paper, we discuss issues arising in applying evolutionary design techniques to a challenging class of problems, namely coordinated collective behaviours of groups of animats. The ultimate aim of our work was to evolve sensory-motor controllers which would give rise to behaviours similar to flocking in birds, schooling in fish, or herding in land animals. These animal behaviours have been extensively studied in the biology literature, and a number of researchers have reported results where animats (both real and simulated) have been constructed to exhibit comparable collective behaviours. Most of the work with animats has involved manually designed controllers, and the indications are that the design of such controllers is a difficult task. Artificial evolution would seem to offer a route by which much of this hard work can be avoided.

In the work we discuss here, we intended to evolve neural-network sensory-motor controllers for virtual “fish” which could move in three dimensions subject to relatively realistic “laws of physics” (i.e. drag, momentum, inertia, etc.). Hence we will talk about “schooling”, but this can be taken to refer also to “flocking”, “swarming”, and “herding”.

While we found it relatively easy to evolve controllers which produced simple collective behaviours such as aggregation or dispersal, we experienced significant difficulties in evolving schooling behaviours. The main aim of this paper is to highlight the problems we encountered, in the belief that they are inherent in the evolutionary approach. We will argue that formulating an effective fitness evaluation function for use in evolving controllers can be at least as difficult as hand-crafting an effective controller design. Although our paper concentrates on

schooling, we believe that this is likely to be a general issue, and can be expected to be experienced over a variety of problem domains.

Section 2 reviews past work in creating animats which exhibit collective behaviours similar to schooling. Section 3 then discusses our evolutionary simulation system. Following this, Section 4 presents results from experiments where a variety of collective behaviours are evolved. Section 5 then discusses our results and the probable reasons for failure to evolve satisfactory schooling. Our conclusions are presented in Section 6.

2 Schooling in Animals and Animats

Fish schooling and other forms of animal aggregations (flocks, herds, and swarms) have always been one of the greatest spectacles that nature can offer. Biologists, zoologists, and lately animat researchers, have shown interest in them. Such natural exhibits are often cited as examples of emergent collective behaviour.

Emergence of collective behaviour is interesting because it offers the possibility of creating complex global behaviour from local interactions between relatively simple agents. It is the sum of these local interactions that makes the system complex as a whole.

It is beyond the scope of this paper to provide a full review of the biology literature: here we briefly discuss hypotheses concerning schooling in fish. We then present a short review of related work in animats.

2.1 Schooling in Fish

Biologists have proposed several different hypotheses for schooling behaviour. Being in a school serves to reduce the risk of being eaten by a predator [PAR82,SHA62,PT86]. It also provides mating efficiency, makes finding food easier, and is a good environment for learning and reducing overall aggression [BILL76]. Some studies have tried to prove that another reason for schooling is energy saving by improved hydrodynamic performance through reducing drag. However, there are contradictory opinions and no conclusive results can be drawn [BILL76,PAR79].

Schooling seems to obey the rules of a distributed model (each individual applying the same set of simple behavioural rules). Each fish takes into account all fish that swim in its neighbourhood, paying more attention to the closest ones and trying to match its velocity and direction with that of its neighbours [PAR82,SHA62]. Fish also try to maintain a constant distance between themselves and their close neighbours.

2.2 Related Animat Work

A seminal work on schooling in animats is Reynolds' "Boids" behavioural animation system [REY87]. In this system, virtual agents exhibit schooling behaviours in 3-D. The behavioural control of each agent can be described by a set of simple rules, but the controller for each agent involves a relatively sophisticated "arbitration" mechanism. Achieving successful schooling in Boids requires fine-tuning a number of parameters in the arbitration mechanism.

Boids used incremental geometric flight to move the agents through the environment. Geometric flight models conservation of momentum and viscous speed damping. Visual perception in each agent is modelled to the extent that it provides the behaviour model with similar information to that available to a "real animal". This information can be seen as being the end result of its perceptual and cognitive processes. However, the generation of such information by a perceptual or cognitive system presents a number of significant difficulties. In this sense, Boids is a "perfect information" system: the relevant variables are continuously available for each agent without noise or error.

Other work on simulation of collective behaviour tends to take a similar approach. Accurate modelling of perception is the weak point of most of them, and various simplifications are usually adopted. Senses like hearing or smell have also been used, but in a similar way as simplified vision [WER92]. Some researchers have modelled additional aspects that affect behaviour, such as internal states (hunger, fear, libido, energy level, fitness, etc.) [TU94A,TU94B,WER92]. Some other pieces of work have treated schools of fish like particle systems, concentrating on attraction and repulsion forces and dynamics [AOK82,NIW94].

Mataric [MAT92] has demonstrated schooling in real robots, working with real sensors. Her research, in common with much other work in collective animat behaviour, relied on hand-crafting the agents' controllers. The schooling behaviour was created by hand-tuning weights which combined the contributions of less complex component behaviours such as aggregation and avoidance. Mataric notes [MAT92,P.438] "Due to the number of tunable parameters involved, *flocking* is the most complex basic interaction implemented in this work so far".

While much work in animat schooling is based on hand-crafting behaviours, there has been some prior work on evolving the controllers for virtual animats in the form of Lisp expressions [REY92], and simple feed-forward neural networks [WER92]. Rucker [RUC93] used an ecosystems model (i.e. multiple interacting animat species) to tune genetically-encoded parameters which

governed a schooling controller inspired by Reynolds' work, but with the arbitration mechanism replaced by a simple linear combination of control variables. Again, all the work on evolving schooling controllers with which we are familiar relies on "perfect information" control variables and hence issues in sensory processing for guidance of schooling are avoided entirely.

2.3 Summary

The following two points are of primary concern to our work:

First, the use of perfect information in virtual animats raises the problem that a sensory system delivering such information is neither biologically plausible nor realistically implementable in real robots. Moreover, deciding which environmental variables are important is often an intuitive or heuristically-guided process, requiring a pre-commitment as to which variables are significant. This pre-commitment to particular variables could be avoided by giving the virtual animats realistic simulated sensory systems and then allowing some adaptation process to determine which factors in the sensory input are relevant for guidance of the desired behaviours.

Second, most of the animat controllers developed to date are hand-coded. The complexity of this manual design task rises as the simulations become more realistic, and the difficulty is presumably most acute in real robot systems such as the one developed by Mataric. Such design processes tend to be difficult, heuristically guided, and time consuming. In principle, it should be possible to use artificial evolution to (semi-) automate this design process. However, the few studies using evolution to design controllers for schooling have worked on simulations with so many simplifying assumptions that their relevance is perhaps questionable.

In principle, it should be possible to use artificial evolution to develop controllers which give rise to schooling behaviours in virtual animats with sensory systems more realistic than those used in prior research. This was the aim of our experiments, described in further detail below.

3 Evolving Animats for Collective Behaviours

The animats in our experiments were loosely modelled on fish, and the aim was to evolve controllers which gave rise to collective behaviours analogous to fish schooling. The primary interest was the notion that coordinated and coherent "global" group behaviours could arise from the interaction of a number of agents,

each of which has access to only "local" information (i.e. that gathered from range-limited sensors). There was no intention to create a faithful simulation of reality such as that developed by Terzopoulos et al. [Tu94A,Tu94B] where the mechanics and hydrodynamics of fishes are taken into account.

Fish are an attractive source of inspiration for two reasons. First, fish move in three spatial dimensions. Avoiding collisions in two-dimensional schooling (e.g. "herding" in terrestrial agents) is more difficult than in the three-dimensional (aquatic or aerial) case where alterations in depth/altitude can be used as an alternative to taking evasive action within the horizontal plane. Second, fish combine visual perception and perception from the "lateral line" pressure sensors. The sensors on the lateral line, which runs longitudinally on each side of the fish, are responsive to local variations in water pressure. Typically, such pressure variations correspond to the fish swimming close to an obstacle (which may be inanimate, such as a rock, or animate, such as another fish), and the indications are that fish use lateral lines to sense both proximity and relative velocity.

The attenuation of light and pressure in water differ, such that the visual system provides distal sensory information while the lateral lines provide proximal information. This mix of distal and proximal sensors is analogous to the use of e.g. visual and tactile sensors on mobile robots.

Our animats had sensory systems which were minimalist approximations to vision and pressure sensors. The physical arrangement of these sensors was fixed during every experiment (not affected by evolution) and lead to a determined architecture for the artificial neural network "controller". Sensory-motor coordination was governed by these artificial neural network "controllers", the weights and thresholds of which were under evolutionary control.

We used neural networks as the basis for the evolving controllers because of the widely accepted argument that their properties of graceful degradation with respect to alterations in weights and thresholds results in smoother evolutionary fitness landscapes. A number of authors have reported successful application of artificial evolution to developing a wide variety of styles of neural network [Kus94].

Further details of the synthetic fish are given in Section 3.1, with details of the evolutionary approach given in Section 3.2. Results from experiments in evolving simple collective behaviours are described in Section 3.3, and results from evolving schooling are discussed in Section 3.4.

3.1 The “Synthetic Fish” Animats

The “fish” animats exist in a 3-D space with bounds on the “depth” (Z) axis but no limits on its horizontal (XY) extent. These bounds are an idealization of fish swimming in the open sea, where they can roam freely but have to stay between the sea-surface and the sea-bed.

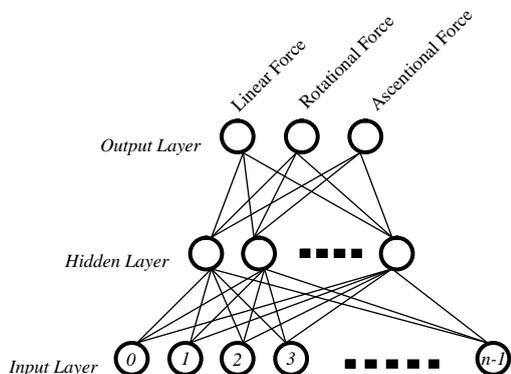


FIGURE 1. Neural network architecture.

The behaviour of each animat is determined by a three-layer feed-forward neural network. Units (i.e. “artificial neurons”) in one layer are fully connected to all units in the next layer, and there are no connections within a layer. The instantaneous output of a unit is a function of the sum of its current inputs. For each unit, the connection weights and the nature of the transfer function were genetically determined: see [ZAE95] for further details. Units in the input layer were either “visual” or “lateral-line” sensors: their inputs values were determined by models of these sensory processes discussed further below.

The neural network for each animat had three output units, which were interpreted as forces which produced: linear acceleration in the Z axis; linear acceleration in the XY plane along the animat’s longitudinal axis; and angular acceleration about the animat’s Z (“yaw”) axis (i.e. altering its orientation in the XY-plane). Simple Newtonian point-kinematics were used to determine the translational and rotational velocities of the animat as the outputs of the network varied: all animats had the same notional masses, moments of inertia, and drag coefficients. The use of drag coefficients gave asymptotic maximum limits on speeds when acceleration forces are applied indefinitely.

For the purposes of simulating sensory input, each animat was considered to have a spherical “body”.

Lateral Lines

Each animat has six “pressure-sensitive” sensors located at the top, bottom, front, back, left, and right of

its spherical body. See Figure 2.

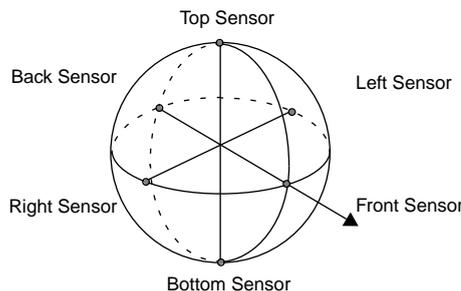


FIGURE 2. Lateral lines diagram.

The vector from the centre of the animat’s body through the location of the sensor forms a normal-vector for the “threshold-plane” of the sensor. The sensor is sensitive to all other bodies (i.e. animats, and the “sea-bed” and “sea-floor” planes) on the distal side of the threshold plane (i.e. the opposite side from the animat’s body). The response from each sensor is a sum of the proximity contributions from these bodies. Proximity contributions from each body are attenuated with distance and with increases in the angle between the threshold-plane normal-vector and the vector from the animat to the nearest point on the body: see [ZAE95] for further details.

Vision

Most biological studies suggest that in fish, the visual system is used to maintain distance and angle to the closest neighbours. Apart from this role, which overlaps with that of the lateral lines, the visual system is used as a long distance perceptual system. In our experiments it could have a very important role at the beginning of the simulation, where depending on the initial conditions animats may be widely scattered, and a long distance sensor is needed to bring them together. In real fish, the visual system is thought to be the primary sensory mechanism underlying the formation of fish schools from several single individuals.

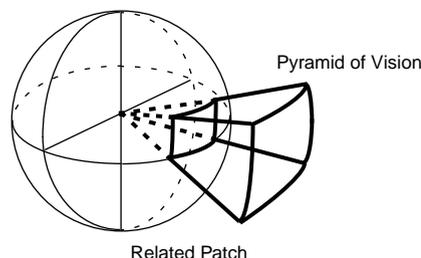


FIGURE 3. Pyramid of vision. The pyramid of vision has infinite volume as it defined by unbounded planes.

Our animats can “see” the world through a primitive 360-degree patched retina. The patches are directly

mapped on to parts of the spherical body. Each patch has a value that depends on how many fish are in the pyramid of vision corresponding to that patch. The pyramid of vision is defined as the part of the space delimited by the segments that go from the centre of the animat through each of the vertices of the patch: see Figure 3.

Statistically, the information produced as the result of the vision system can be understood as a histogram of the surrounding environment. Every column of the histogram represents one patch of the visual system. The value of the column stands for the activation of that patch. The activation of a given patch is proportional to the number of fish contained in the associated pyramid of vision and how far away they are.

The number of patches was not under genetic control. In all our experiments there were nine patches in a 3x3 grid on the frontal hemisphere of the animat. We tried experiments where the rear hemisphere was a single patch, and also where it was a 2x2 grid of patches.

Having more patches (i.e. higher “resolution”) at the front of the animat is an attempt to give it more biological plausibility. This approach models an approximation to foveal vision. Whether this potential advantage is exploited by the synthetic fish or not depends entirely on the evolutionary process.

3.2 The Evolutionary Process

In our experiments, the genotypes that the genetic algorithm manipulates encode the arbitration systems of the synthetic fish. Every genotype encodes the parameters for the transfer function of every unit of the neural net controller as well as the weights between connected units. The organisation of the genotype is very simple. Every gene is coded as a double precision floating point number in the range [0.0,1.0). The sequence in which the units are placed in the genotype is also straightforward. Starting from the first unit of the input layer, following with the units from the hidden layer and finally the output units: see Figure 4.

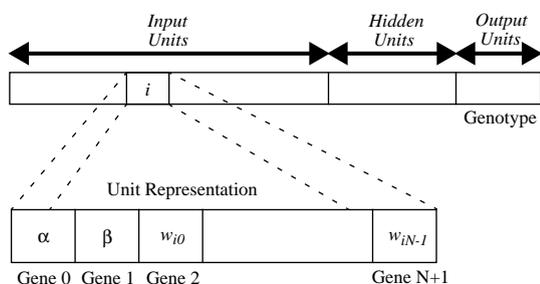


FIGURE 4. Genotype structure: α and β are parameters for the unit’s transfer function.

The genetic operators used were one point crossover and mutation. Both operators act at the gene level. Taking into account that most of the genes encode neural network weights, mutation acts in an incremental fashion rather than as a complete random operator. In practice, it works by adding a small (compared to the gene value range) random value to the previous value of the gene.

A typical experiment, with nine patches at the front and one big patch at the back, gave a neural network configuration of 16 input units, 5 hidden units and 3 output units. Using our genetic encoding this configuration amounts to a genotype length of 143. To set the mutation rate and crossover rate of the genetic algorithm we used information from the convergence measures (average and standard deviations of the mean Euclidean distance to the elite and mean genotypes). We found it appropriate to set the mutation rate at approximately one gene mutated per genotype in the next population. Selection was rank-based, with a near-quadratic exponential reduction in probability of an animat being chosen to reproduce.

As we were interested in collective behaviours from groups of homogeneous animats, we evaluated an individual genotype by monitoring the behaviour of groups of clones of that genotype. The genotype was used to set the parameters for the neural networks of a number of animats. Simulation of these animats moving in the 3D environment was then monitored to determine the fitness of the genotype.

The genetic algorithm was usually run for 100 generations. We did most of our simulations with 4 or 5 animats and using only a single set of initial conditions (only one evaluation per genotype). Each simulation was run for 2000 steps.

We also tried some experiments using different settings: more animats, more sets of initial conditions, more simulated time, more generations, etc. None of these variations significantly affected the qualitative nature of the results presented below.

3.3 Dispersal and Aggregation: Success

In order to confirm that our evolutionary system could operate successfully (i.e. that we had the various parameters such as mutation rate, crossover rate, etc. set correctly), we first ran experiments where collective behaviours less challenging than schooling were evolved. We evolved controllers for “dispersal” and “aggregation” behaviours. Both of these require that the animats be sensitive to the actions of others in the group. Indeed, the only significant bodies in the environment with which an individual can interact are other animats.

Dispersal

The evaluation function was based on maximising the distance to the closest neighbour of every fish. This behaviour was the easiest to evolve: a rule such as, “swim to where you can detect nobody” didn’t seem very difficult to hard-wire in the neural network architecture. Results from a typical dispersal experiment are shown in Figure 5.

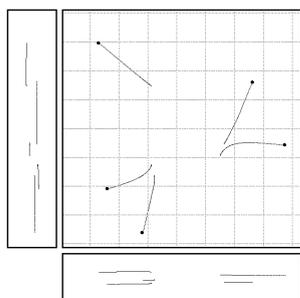


FIGURE 5. This figure shows the first 100 time steps of the dispersal experiment. Central view is an XY projection (top view), while left and bottom are YZ and XZ lateral projections respectively. When the trial starts, the animats orient away from each other and then scatter in straight line-paths until the end of the trial.

Aggregation

The aggregation evaluation function was based on minimising the distance to the closest neighbours. The optimal solution would be to make all the animats collide in a single point in space and not move any more. Our experiments demonstrated that because the animats had realistic dynamics (inertia, maximum turning radius, etc.), the optimal solution would not happen. Instead, the animats stayed moving in a fairly stable bounded region of the space. This result can be considered fairly good because it shows certain degree of evolutionary adaptation to the given task, and that there is a defined perception-reaction function coded in the neural network (we tested that by disconnecting the sensory inputs from the neural network and observing the resulting behaviour, which was similar to a random scattering behaviour). Figure 6 shows results from a typical aggregation experiment.

As we’ve mentioned before, for all these experiments, the space in which fish could swim was unbounded in the X-Y plane. We performed similar experiments with bounded environments (upright cylinders of different sizes). For these we added to the lateral lines sensory system the capability of detecting the proximity of the bounds of the arena as if they were any other object or fish in the simulation.

We discovered that the genetic algorithm found a way of maximising some of our evaluation functions based on the fact that the simulation arena was bounded.

An example of this, is a primitive aggregation behaviour that we observed. Even though it looked like emergent collective behaviour, we discovered that it was due to the individual interactions of each fish with the environment (bounds of the arena) rather than with the other animats.

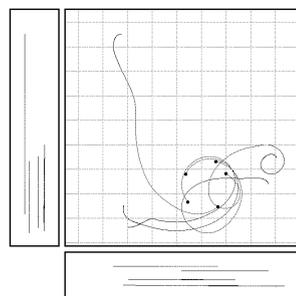


FIGURE 6. This figure shows the traces of the animats over the first 400 time-steps. From initial random positions, the animats converge to occupy a bounded region of space. Once aggregated, they all move on approximately concentric circular paths, as if they were “chasing” one another: the start of this can be seen here.

3.4 Schooling: Failure?

To evolve schooling behaviour we had to provide an evaluation function that rewarded individuals that schooled. This seems obvious, but we were unable to find a scientific quantitative measure in the literature. Most papers suggest that schooling is the property of moving in the same direction, at the same speed and at a preferred distance from your nearby neighbours. Following those principles we formulated a variety of evaluation functions which we believed would reward appropriate behaviours.

Most of our evaluation functions involved maximising a sum of values from sub-functions which rewarded the component behaviours which other authors have proposed as combining to form schooling. The sub-functions gave rewards for “maintaining the preferred distance to neighbours” or “travelling at the same speed as neighbours”, and so on. The individual rewarding functions were basically Gaussian functions that returned values proportional to the extent that the desired schooling quality was accomplished.

The wide variety of behaviours that we obtained ranged from some fairly random behaviours, through curious fixed behaviours such as every fish swimming in small circles (due to constant outputs in the neural network), to the one shown in Figure 7, where a form of schooling has actually evolved. This example was rewarded highly by our evaluation function, but it lacks realism: the two schools swim in circles. Unfortunately, this kind of primitive schooling was far from what the average person would acknowledge as “realistic” schooling.

The biological literature indicates that there are factors (migratory urges, gradients in temperature or light, etc.) that make animals school in a certain direction rather than in any other. In our very impoverished environment, where the only other thing that you can detect is other animats, there is no reason to move in any particular direction. Providing some biologically plausible environment properties that the animats could detect (via appropriate sensors) might make it easier to evolve realistic schooling. But there are other possible reasons for our lack of success, discussed in the next section.

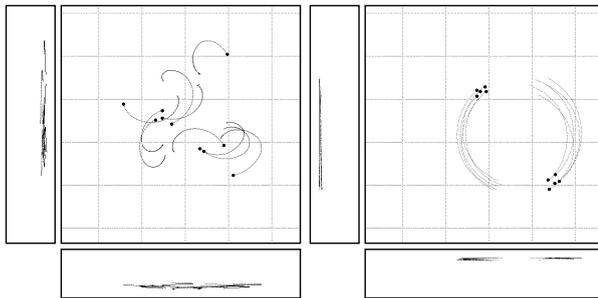


FIGURE 7. The left-hand figure shows the starting positions and the first 100 time steps of a schooling simulation. The genotype producing this behaviour was the elite from the final generation. Dots represent the positions of the animats at time step 100. The right-hand figure shows the same experiment after 1800 simulated time steps. Only the traces for the last 100 time steps are shown. Here we can see that the animats have formed two schools and swim side by side in circles.

4 Discussion: What went wrong?

Our results seem to indicate that our approach is only capable of generating very simple collective behaviours. We therefore had to consider the possibility that some part or parts of our evolutionary approach were unsuitable to evolve more complex behaviours. The main possibilities that concerned us were: That the neural network architectures, and the nature of the sensory system, were inappropriate for generating schooling behaviour; and that the evaluation function was wrong.

4.1 The Controller Architecture

We had a strong suspicion that a feed-forward neural net architecture like ours is not sufficient to process the kind of information that the sensory system produces. The sensory information is effectively a “snapshot” of the surrounding environment at a given time. From the point of view of an individual animat, there is no link between one snapshot and the next. The perception-reaction function is only based on instantaneous information. Differential information, such as rates of approach and

directions to and of other objects, might be essential for an animat to show schooling behaviours. To calculate this kind of information, a recurrent neural network rather than a feed-forward one might be needed. Alternatively, the current sensors could be augmented with new sensors which provide differential input, such as “image motion” in the visual system.

For this reason, we replaced the sensory systems described above with a “perfect information” system similar to that used in the flocking software animats described in Section 2.2. Now the inputs to the networks were values such as the nearest neighbour’s rate of approach and distance. None of these experiments gave results significantly better than the earlier experiments.

4.2 The Evaluation Function

In order to check our evaluation function, we implemented hand-crafted animat controllers based on rules argued to be capable of generating schooling behaviour. We developed a prototype partly inspired by Rucker’s ideas [RUC93] but with some conceptual modifications needed to adapt it to our more realistic kinematics model. The appeal of this prototype is that it is a simple system, which appears to produce behaviours comparable to Reynolds’ more complex Boids controller. It uses a mechanism where the response of an individual animat is based on a weighted combination of a small number of 3D vectors. Alterations to the animat’s tangent vector (i.e. its direction of movement) are based on two other vectors: the nearest neighbour’s tangent vector and the vector from the animat to a centroid point. Alterations in speed are determined solely by the distance to the nearest neighbour.

We fine-tuned all the parameters of our version of the Rucker controller (which was a time-consuming, purely heuristically-guided process), until the behaviour of the animats gave what we (subjectively) judged to be realistic schooling. We then analysed the performance of our evaluation function by examining the contributions of the sub-functions to the final fitness measure, and by checking that this good schooling behaviour scored close to the theoretical maximum score. The scores of some sub-optimal controllers were also monitored, to estimate the smoothness of the evaluation function. While this allows us to check that a given evaluation function rewards appropriate behaviour and gives intermediate scores to intermediate behaviours, there is unfortunately no method we can use to make sure that the evaluation function is not going to reward something that is not schooling.

One significant issue is that, in contrast to evolving controllers for solitary agents, we need to develop evalu-

ation functions which give a global measure of schooling from the observed behaviour of a group of individuals. Although schooling is defined as a collective phenomenon, monitoring schooling is fundamentally based on monitoring the behaviour of individuals relative to the rest of the group. Thus, there are at least four important decisions to be made in developing an evaluation function. The first concerns which individuals in the group are monitored; the second concerns which variables are monitored for those individuals; the third concerns how those variables are combined to give some instantaneous measure of group behaviour; and the fourth concerns how the instantaneous measure is integrated over the duration of the trial to arrive at the final global measure of group performance. At each of these decision stages there are potentially many possible alternative approaches, and often there is nothing more than heuristic guidance as to which one to choose. So the space of possible evaluation functions is vast.

Michalewicz [Mic94] makes the distinction between multimodal optimisation problems (with more than one optimum) and multiobjective problems (optimise more than one thing). Our problem basically belongs to both categories.

To form a school from a scattered group of animats, the group must first aggregate and then orient to the same direction. Thus the evolution of schooling is multi-objective in that the initial generations are likely to be selected primarily for their ability to aggregate. There is a danger that this behaviour dominates over the others. We have observed this phenomenon in some of our evaluation functions.

We performed some experiments where we used dynamically incremental evaluation functions, where additional sub-functions were activated at different stages in an individual trial only when certain conditions were reached. For example, while the group were in the process of aggregating, measures of aggregation were weighted more heavily in the overall evaluation function than were those for schooling; once aggregated, schooling was given more significance than aggregation. Again, finding an appropriate balance between the different contributions in the evaluation function proved to be the limiting factor.

5 From a Significant Example to a General Principle?

If a standard GA is to be used to evolve schooling behaviours in animats, the evaluation function must be some quantitative measure of schooling. To the best of our knowledge, nowhere in the literature on collective

behaviours (in either animals or animats) is such a quantitative measure employed: the observed behaviours are described as schooling, flocking, or herding on the basis of appeal to intuitive and subjective notions of these behaviours.

It could be argued that the controllers for Mataric's robots or Reynolds' Boids constitute implicit definitions: i.e., if a group of agents do what the robots or boids do, then the group is schooling. Even if an operational definition was available which allows for some spatio-temporal pattern of activity to be classified as "schooling" or "not-schooling", this would not be sufficient for the needs of an evolutionary approach: to apply a standard GA, a quantitative function is required which gives a graded response that in some way reflects the *degree* to which the observed behaviour can be classified as schooling.

Formulating a quantitative function is not difficult: so long as the function yields a reasonably smooth fitness landscape, the GA is likely to find controller architectures which are significant improvements on the initial random designs. But formulating an *effective* function can be very hard: even when we first checked our evaluation functions against the Rucker-based controllers, the indications from subsequent evolutionary experiments were that many functions which gave appropriate rewards to schooling also gave rewards to manifestly degenerate behaviours: the space of possible behaviours satisfying the functions was not sufficiently constrained. Introducing extra constraints often appeared to result in over-severe fitness landscapes, where improvements on initial random designs rarely occurred.

Of course, the fitness landscape in a particular artificial evolution experiment is not solely determined by the evaluation function: it is the result of an interaction between the evaluation function and other factors such as the genetic encoding (which determines the space of possible controller architectures) and the agents' sensory-motor interactions with the environment(s) (which determine the space of possible behaviours, the ultimate "phenotypes" in the system). Nevertheless, the fact that we could successfully evolve controllers which produced simpler collective behaviours such as aggregation and dispersal indicates that it was reasonable to expect schooling to also be evolvable.

It could reasonably be argued that real schooling behaviours in real animals arise because of the complex interaction of a number of factors, and that our approach failed because the simulations lacked sufficient complexity. We have much sympathy for this argument, and other authors have already demonstrated that more complex evolutionary simulations can show interesting results. Two promising developments in this direction are Rey-

nolds' [REY92] use of a hard-wired "predator" animat to select for controllers giving coordinated group motion in "prey" animats, and Rucker's [Ruc93] use of interactions within an "ecosystem" to genetically tune parameters for his Boids-like animat controllers. However, in the absence of a schooling metric, it is difficult to judge to what extent the work of either of these authors can be regarded as the evolution of schooling: neither author attempted to quantify the degree to which their evolved animats were exhibiting schooling behaviours.

So perhaps our approach failed to capture the complexity necessary for the successful evolution of schooling. If, in order to evolve controllers for a behaviour as intuitive and simply stated as schooling, it is necessary to construct a complex ecosystem, then so be it. But, even with a more complex system, the absence of a quantitative measure of schooling implies that deciding whether schooling is occurring or not becomes a task that requires *manual* monitoring of the evolutionary process. So although no human labour is involved in designing the controllers, human labour will be needed (perhaps for days or weeks) to watch out for the possibility that the evolving animats have started to school.

Furthermore, if our lack of success was a consequence of lack of complexity, then this has significant negative implications for arguments advocating artificial evolution as a labour-saving alternative to manual design. The hand-crafted controllers in Reynolds' Boids and Mataric's robots are the result of skill and creativity on the part of their designers. In principle, or so the story goes, the need for skill and creativity can be reduced by using artificial evolution: it is necessary merely to formulate an evaluation function and apply a GA. The problem our work highlights is that developing an appropriate evaluation function can be a difficult, time-consuming, heuristically-guided process requiring skill and creativity too.

This is compounded by the fact that, to test an evaluation function, it is generally necessary to perform more than one evolutionary experiment with that function. As soon as the function can be demonstrated to produce satisfactory results, it can be declared a success. In the absence of a success, however, it is generally necessary to perform more than one experiment: in principle one should continue generating failures from independent experiments until statistically significant conclusions can be drawn, and each experiment should be run for long enough to give a realistic chance for successful evolution to occur. Of course, this brute force approach can be avoided if analysis of the outcome of the failed experiments reveals *why* the evaluation function is inappropriate, but such analysis can be a difficult and time-consuming task. The point here is that this "meta-evalua-

tion" process (i.e. evaluating the effectiveness of the evaluation function) can take a lot of time and effort, perhaps more so than if a manual-design approach had been taken. Gauging the effectiveness of a particular evaluation function can be a very expensive process. The costs of evolving a design can be very high, because many intermediate solutions will be evaluated through the evolutionary search process. Given that a design is already partially specified in the evaluation function, large amounts of time may be saved by hand-crafting possible designs and then judging their behaviours.

In some cases, the evaluation function becomes sufficiently complex that it is little less than a full specification for the desired behaviour. Even if it is an effective evaluation function, there is an implication that the desired behaviour must be understood in advance, and so much of the design work is done before the GA starts.

Additionally, it is often necessary to employ heuristic procedures to fine-tune the evaluation function. Conceptually, it seems hard to distinguish this from the process of heuristically fine-tuning a manual design. In practice, the former is far more difficult and time consuming because the effects of the changes are only seen after the evolutionary process has run. It may also be the case that the consequences of changes in a hand-crafted design are more easily understood and lead to a faster design cycle than changes in a evaluation function.

6 Conclusions

Our overall conclusion is that, for schooling behaviours at least, the time and effort taken to develop an evolutionary system is probably more than the time taken to develop a similar controller using manual design techniques. This is probably due to the fact that schooling is a sophisticated, complex, and incompletely understood phenomenon, which until now has been subjectively defined.

This paper has concentrated on what are essentially negative results. But negative results are still results. Our primary message is one of caution: our reasons for failing to successfully evolve controllers analogous to those hand-crafted by other authors are, we believe, rooted in the difficulties of formulating an effective evaluation function. This aspect of the application of artificial evolution to the development of sensory-motor controllers for autonomous agents is currently much more of an art than a science; guided as it is by heuristics and time-consuming trial-and-error techniques. Given the relative recency of the development and application of these techniques, it is perhaps no surprise to find such an important aspect of the field at a pre-theoretical stage.

Interest in artificial evolution as a replacement to manual engineering-design techniques is likely to increase significantly when the evolved systems require less human effort to create than would have been the case if they had been designed by hand. For this to happen, we anticipate, it will be necessary for the methodology to be advanced beyond the pre-theoretical stage, to the point where there are “predictive engineering” practices that can be deployed to (help) formulate the evaluation function in much the same way that an architectural engineer can (hopefully) predict whether a certain arrangement of walls will support a certain style of roof. How such practices might be developed is, at present, an open question.

Acknowledgements

Initial work on this paper was done while financial support for NZ came from GRANT DETAILS. DC thanks Inman Harvey for discussions. JB ACKS?

References

- [AOK82] Aoki I. (1982) “A simulation study on the schooling mechanism in fish”. *Bulletin of the Japanese Society of Scientific Fisheries* 48. Pages 1081-1088.
- [BIL76] Bill R.G. and Hermkind W.F. (1976) “Drag reduction by formation in spiny lobsters”. *Science* 193. Pages 1146-1148
- [KUS94] Kuscu I. and Thornton C. (1994) “Design of artificial neural networks using genetic algorithms”. Review and prospect in *Proceedings of Third Turkish Symposium on Artificial Intelligence and Neural Networks*, Bozsahin, C. et al. editors. Pages 411-420.
- [MAT92] Mataric M.J. (1992) “Designing Emergent Behaviors: From Local Interactions to Collective Intelligence”. in *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior (SAB92)*, Meyer, Roitblat and Wilson editors, MIT Press, Cambridge, Massachusetts. Pages 432-441
- [MAT95] Mataric M.J. and Cliff D. (1995) “Challenges In Evolving Controllers for Physical Robots”, submitted to “Evolutional Robotics”, special issue of *Robotics and Autonomous Systems*, also Brandeis University Computer Science Technical Report CS-95-184.
- [MIC94] Michalewicz Z. (1994) *Genetic Algorithms + Data Structures = Evolution Programs*. Second Edition. Springer Verlag, New York 1994.
- [NIW94] Niwa H.S. (1994) “Self-organizing Dynamic Model of Fish Schooling”. *Journal of Theoretical Biology* 171. Pages 123-136
- [PAR79] Partridge B.L. and Pitcher T.J. (1979) “Evidence against a hydrodynamic function of fish schools”. *Nature* 279. Pages 418-419
- [PAR82] Partridge B.L. (1982) “The structure and function of fish schools”. *Scientific American* 246. Pages 90-99
- [PIT86] Pitcher T.J. (1986) “The functions of shoaling behaviour”. in Pitcher T.J. (eds.) *The Behaviour of Teleost Fishes*. Croom Helm, London. Pages 294-337
- [REY87] Reynolds C.W. (1987) “Flocks, Herds, and Schools: A Distributed Behavioral Model”. *ACM Computer Graphics*, 21(4) (SIGGRAPH '87 Conference Proceedings). Pages 25-34
- [REY92] Reynolds C.W. (1992) “An Evolved, Vision-Based Behavioral Model of Coordinated Group Motion”. in *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior (SAB92)*, Meyer, Roitblat and Wilson editors, MIT Press, Cambridge, Massachusetts. Pages 384-392
- [RUC93] Rucker R. (1993) *Artificial Life Lab*. The Waite Group Press. Corte Madera. CA.
- [SHA62] Shaw E. (1962) “The schooling of fishes”. *Scientific American* 206: Pages 128-138
- [TU94A] Tu X. and Terzopoulos D. (1994) “Artificial Fishes: Physics, Locomotion, Perception, Behavior”. *Proceedings of ACM SIGGRAPH'94*, Orlando, FL, July, 1994, in *ACM Computer Graphics Proceedings*. Pages 43-50
- [TU94B] Terzopoulos D., Tu X. and Grzeszczuk R. (1994) “Artificial Fishes with Autonomous Locomotion, Perception, Behavior, and Learning in a Simulated Physical World”. in *Artificial Life IV: Proc. of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, Cambridge, MA. Pages 17-27
- [WER92] Werner G.M. and Dyer M.G. (1992) “Evolution of Herding Behavior in Artificial Animals”. in *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior (SAB92)*, Meyer, Roitblat and Wilson editors, MIT Press, Cambridge, Massachusetts. Pages 384-392
- [ZAE95] Zaera N. (1995) “Simulated Schooling: Collective Behaviours in Synthetic Fish”. *Unpublished MSc Thesis*. School of Cognitive and Computing Sciences, University of Sussex, Brighton.