

# Evolving Collective Behavior in an Artificial Ecology

---

**Christopher R. Ward**

University of Nottingham  
School of Computer Science

**Fernand Gobet\***

University of Nottingham  
School of Psychology

**Graham Kendall**

University of Nottingham  
School of Computer Science

**Abstract** Collective behavior refers to coordinated group motion, common to many animals. The dynamics of a group can be seen as a distributed model, each “animal” applying the same rule set. This study investigates the use of evolved sensory controllers to produce schooling behavior. A set of artificial creatures “live” in an artificial world with hazards and food. Each creature has a simple artificial neural network brain that controls movement in different situations. A chromosome encodes the network structure and weights, which may be combined using artificial evolution with another chromosome, if a creature should choose to mate. Prey and predators coevolve without an explicit fitness function for schooling to produce sophisticated, nondeterministic, behavior. The work highlights the role of species’ physiology in understanding behavior and the role of the environment in encouraging the development of sensory systems.

---

## **Keywords**

ALife, artificial ecology, coexistence of prey and predator, evolution, fish, flocking, schooling, sensor

---

## **I Introduction**

Collective behaviors such as conformity, obedience, and leadership have been studied for years in social psychology. Yet we can see primitive characteristics of these behaviors in animals, such as fish schools, insect swarms, and bird flocks. Even across such a diverse set of species, the dynamics of the group are very similar. Partridge [18] states that collective behavior occurs when animals move “in unison, more like a single organism than a collection of individuals” (p. 90). Movement is dependent on the characteristics of the animal; for example insects can fly in three dimensions, unlike sheep, which can only move in two dimensions. The circumstances that stimulate movement differ too, for example, the presence of prey or cold climates. In this article we consider the role of sensory configuration as an important aspect in determining collective behavior.

From an evolutionary standpoint, it is understandable why collective behaviors are prevalent in many species. Biologists propose several hypotheses for collective behavior, including reducing the risk of being eaten by a predator, providing mating efficiency, making the search for food easier, and providing a good environment for learning and reducing overall aggression [23]. Collective motion may also save energy through reducing drag [17].

Zoologists and other scientists have studied collective behavior in nature for a long time, but these phenomena have proven difficult to study objectively without threatening ecological validity. Partridge [18], for example, studied collective behavior in fish

---

\* Correspondence to: F. Gobet, University of Nottingham, School of Psychology, Nottingham NG7 2RD, UK, frg@psychology.nottingham.ac.uk; www.cs.nott.ac.uk/~gjk/alife/ward2001/

using a circular tank 10 m in diameter with a central gantry. Experimenters projected a light over the fish, which were conditioned to be attracted to it. Fish were marked with a number and, as the gantry moved, the fish were tracked by measuring various distances to investigate the adjustment of position within the school.

In recent years, computer modeling and simulation have provided a concrete way to test and derive new theories of collective behavior. Reynolds [21] presented the seminal artificial life (ALife) work in collective behavior. His program BOIDS, which implemented artificial birds, did not make any pretense to represent the behavior of birds. Instead, its objective was to produce convincing flocking behavior. Each BOID executed three simple rules or tendencies in the presence of neighbors. Using these three simple rules, complex global behavior emerged from simple local interactions. The results have been reproduced many times. Some variation on the initial rule set and the method of obtaining neighbors has been explored. In general, changing depth of vision and parameters such as tendency to change velocity and heading in response to neighbors can drastically alter the structure of the flock.

Since Reynolds [21], little research has been carried out to devise a rule set that can produce more realistic behavior without compromising the sheer simplicity of the original work. Mataric [14] successfully developed robots to produce flocking behavior. Mataric states that collective behavior is the weighted combination of a number of basic interactions: collision avoidance, following, dispersion, aggregation, and homing. By programming each of these behaviors into several robots and then setting a weight that determined which was more likely to execute, Mataric was able to produce some fairly sophisticated collective behavior.

These “behavioral” models are suitable for defining what characteristics to look for when identifying collective behavior. However, they were “hand written.” Rather than consider the environment and physiology of the species, they are based on some concept about what principles might be considered important [22]. By focusing on behavior alone, these models create a deterministic, one-dimensional controller. Knowledge is generally represented in productions (i.e., “if . . . then . . .” rules) and these models frequently neglect sensory modalities. Here we describe the design of a realistic controller called E-BOIDS capable of a variety of behaviors, where schooling is a truly emergent property of the system, rather than deterministically stated. E-BOIDS is an acronym for evolving BOIDS, in recognition of Reynolds’ work.

The article is organized as follows. Section 2 briefly discusses the existing literature on collective behavior and fish physiology. Section 3 outlines the rather unique features of the implementation. Section 4 describes the design of our system. Section 5 presents a thorough statistical analysis of the results. Finally, we discuss some conclusions.

## 2 Related Work

Husbands, Harvey, Cliff, and Miller [13] claim that there is no evidence that humans are capable of designing systems exhibiting emergent behavior involving many interactions between constituent parts. They suggest the “use of artificial evolution to fully, or partially, automate the design process” (p. 133). With this in mind we can briefly review the work that has attempted to use evolutionary methodologies [20, 22, 23, 24]. Each has had mixed success, with typical problems in assigning fitness to the population. Zaera, Cliff, and Bruten [24] applied the BOIDS rules to synthetic fish, but these rules proved to be insufficient in describing all the characteristics of a school. In fact the authors state that “to the best of our knowledge, nowhere in the literature on collective behaviors (in either animals or animats) is a quantitative scalar measure of schooling employed” (p. 642). A more dynamic measure of fitness, such as that used in PolyWorld [20], is an energy score in combination with an artificial ecology. This

encourages a diverse range of creatures, without discriminating against other interesting behaviors. The important consideration here is to consider the biology of the species (e.g., the sensory modalities available) and the influence of the environment on its evolution.

Fortunately, there is no shortage of literature on every aspect of fish behavior and biology. Some of this background must be explored in order to design the artificial sensory systems necessary to evolve schooling behavior. First of all, let us clarify the terms “schools” and “shoals,” since they are often used interchangeably. Shoaling is where a group of fish occupy an area or territory, whereas schooling is where a group of fish act as a single organism. A school of fish is a fascinating display of nature, hundreds of small fish gliding in unison, more like a single organism than a collection of individuals. There is a highly coordinated structure, yet there is no leader or external stimulus that prompts the form of this polarized configuration. It is estimated that “out of 20,000 species of fish, more than 10,000 species collect in schools during some part of their lives” [18, p. 91], which suggests a common evolutionary history.

Schooling would appear to obey the rules of a distributed model (each individual applying the same set of simple behavioral rules). Each fish takes into account all fish that swim in its neighborhood, paying more attention to the closest ones and trying to match its velocity and direction with that of its neighbors [18]. Yet, it is known from Zaera, Cliff, and Bruten’s [24] work that the BOIDS rules do not describe all these properties, since the rules proved insufficient as a fitness function. And further supporting this notion, schooling varies from one species to another. It is known that some fish spend all or most of their time in a school (obligate schoolers) whereas some join occasionally (facultative schoolers). A slightly more sophisticated model is needed, and for this a flexible representation of the fish’s knowledge (i.e., its sensory systems) must be designed.

### 3 Modeling E-BOIDS

This section outlines the components of our system. First, we consider the flexible controller (i.e., the neural network) on which we can build a host of behaviors. Second, we take a look at distributing the controllers to creatures in the environment using chromosomes and evolving the system using genetic algorithms (GAs) such that profitable behaviors are retained in later generations. Finally, we consider the environment in which the creatures “live.”

#### 3.1 Artificial Neural Network—Fish Biology

Fish have a remarkably small brain, yet they are known to produce some of the most complex collective behavior in the animal kingdom. As there is little evidence to suggest that schooling is voluntary, a reactive model could be a suitable methodology. Braitenberg [1] explored the use of reactive robots in a set of experiments that looked at the potential of building synthetic creatures using a set of sensors. It was found that by varying the connections between sensors and motor actions on a simple base unit a whole host of behaviors could be observed. Many ALife researchers use this kind of methodology to design robotic creatures, or hardware animats. As discussed, the design of the architecture can be troublesome if we want to achieve complex behavior, especially if it involves changing physically wired connections.

An artificial neural network (ANN) can be applied using the same methodology, but because it is software based, the structure can be changed at will. The ANN must be designed to facilitate schooling and other behaviors such as predation. If there is not enough freedom in the architecture, the range of behaviors will be limited. This question will be taken up in the next section.

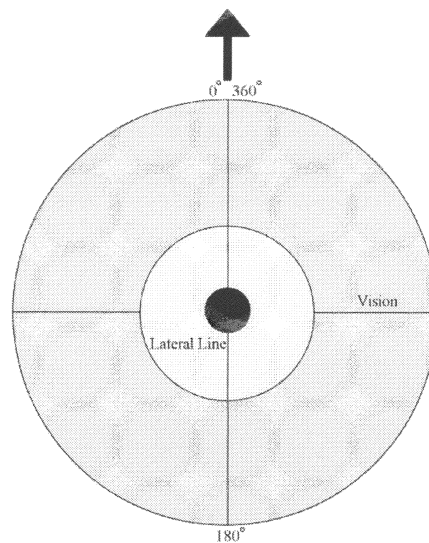


Figure 1. An E-BOID projects senses across its environment.

### 3.1.1 Input Units—Fish Senses

The different sensory modalities can be modeled to provide the organisms with different types of information about their environment. The input units need to be directional and able to distinguish between objects. For example, on seeing a predator on the left, the prey learns to turn right. To achieve directionality, a “pyramid of vision system” is used (shown in Figure 1). The neighborhood region for each sense can be divided into sectors. For example, if there were two sectors, one would capture all objects between 0 and 180° and the other between 180 and 360°. The depth and the number of these sectors differ among sensory types and organisms. The resolution of these patches could prove to be very important.

Given that we are interested in the essence of the creature, only two of the most important senses are considered in this article. However, it should be kept in mind that fish possess other senses, such as chemical releases. We now describe the two senses in question, lateral line and vision.

**Lateral line** Zaera, Cliff, and Bruten [24] represented two major fish senses, the lateral line and vision. The lateral line is the hearing system. Sound vibrations are transmitted from the water through the fish’s body to its internal ears. The ear canal contains sensory hair cells that respond to the movement of the liquid within the fish’s ear, just like the human inner ear. The lateral line is used to detect local changes in water pressure within one or two body lengths of the fish [18]. It acts as a repulsion force, ensuring that fish maintain their personal space.

**Vision** Although water permits only limited visibility, fish are capable of distinguishing predators from prey using vision [18]. However, when facing a school, predators are confused by a sensory overload, unable to select any one member. The visual system can see far beyond the lateral line range. It is used to maintain distance and angle to the closest neighbors [23]. Schooling fish compare information from the eyes with the information from the lateral line [18]. Combining these two into a single ANN classifier is a suitable model for implementing this behavior.

In PolyWorld [20], the investigators rendered the viewpoint of each creature, which was passed into the ANN brain. This is actually extremely complicated, because we must also encompass motion and depth information. It also violates the principle

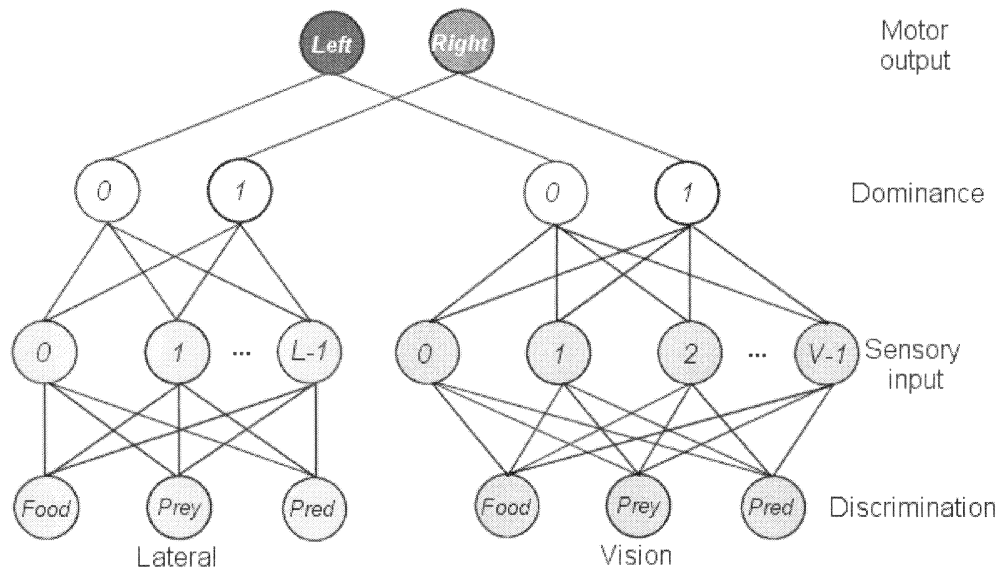


Figure 2. Network architecture of an E-BOID, showing the lateral line input units (where  $L$  is maximum) and the vision input units (where  $V$  is maximum). Each input unit refers to a sensory sector; for example, 0 might refer to  $0-45^\circ$ . For each input unit the discrimination weights are applied for that type of input (i.e., lateral or vision) for food, prey, and predators. All these inputs feed forward in a fully connected manner to the intermediate layer (dominance). Intermediate nodes modulate the balance between lateral and vision sensor inputs (dominance). Their output is passed to the corresponding output unit that determines the motor output (turn left or right).

of ALife, which is to capture the essence of the system. We shall assume that such information is passed directly into the ANN (i.e., a perfect sensory system). This is achieved by a function of both the area occupied by objects and their distance away from the center of the creature.

### 3.1.2 Output Units—Fish Behavior

Werner and Dyer [23] and PolyWorld [20] represented a different behavior in each of the output nodes. This study takes a simple approach. Instead of discrete output values (i.e., 1 or 0) that limit a creature to left/right or mate/not mate, nodes take continuous values to create more believable motion, capable of behaviors that are less limiting.

Werner and Dyer represented three output values, in their three-dimensional system. We propose just two for this two-dimensional system, a turn left and turn right node. If the input to the left node is 1 (the maximum value) then that would result in a  $180^\circ$  turn counterclockwise, an input of the same value to the right node would result in a  $180^\circ$  turn clockwise. Note that this change is always with respect to the current angle; therefore if both output nodes were excited equally, the organism would not change direction (a useful property).

It is assumed that a creature is always moving, and that this motion is always constant. Therefore an organism cannot suddenly speed up or slow down. The speed of a creature is determined by the body size, so a larger body will swim faster than a smaller one (to derive the coordinates of the movement, the tail is moved to where the head was previously).

### 3.1.3 Network Architecture

The ANN is a feed-forward, fully connected network (see Figure 2). A weighting of 0 will mean the connection does not pass its input value. An input node represents a

patch of sensory input that projects a pyramid of “vision” that identifies the position (with respect to itself) and the type of organism in the neighborhood. The attention given to each type of organism when in the different types of sensory modalities is scored between 0 and 1. This means a creature might, for example, pay more attention to a fellow prey organism in the visual system, but be less interested if it is in the lateral range (to avoid collision, for example). These discriminatory weightings are stored independently from the ANN.

The network output is computed as follows. First, the contribution of each sensory modality and sector is calculated. Each sector corresponds to one input node, and each input node sends a signal to the motor output layer (turn left and right) through the dominance nodes (see Figure 2). There are three recognized objects in the world (food, prey and predator), and each creature has weights that dictate the attention paid to these objects (discrimination nodes).

The objects in the environment are translated into sensory signals as a function of both the area occupied by objects and their distance away from the center of the creature. The contribution of each object to a sector is calculated by multiplying its input activation by its discrimination weight (the discrimination weights are the same for all the sectors within a sensory modality). With the contributions of the input nodes established, the input is fed forward additively in a fully connected manner to the dominance layer (see Figure 2). This same calculation is applied for both lateral and visual sector inputs. The weights of the dominance layer determine the balance between the lateral and visual systems. For each sensory system, the inputs to left and right nodes are simply multiplied by the corresponding dominance weight.

The same base network structure is shared between all creatures in the environment, which allows them to be easily combined (i.e., mated). Independent weights are stored for each creature. We can consider the implications of these sensory weights on behavior in a simple example. A good predator sensory system might turn left if it sees prey on the left, and turn right if it sees prey on the right. To achieve this, we have weights: left sense to left motor = 1, left sense to right motor = 0, right sense to right motor = 1, and right sense to left motor = 0. This sensory network would not be of much use for prey. In contrast, prey would be better served by a network that turned themselves away from predators (the inverse of the predator network). Although this is a trivial example, we would not wish to have to hand code such values, because the emergence of interactions between different weightings and between behaviors of different creatures cannot be foreseen. The next section looks at techniques to enable the evolution of these values.

### 3.2 Genetic Algorithms

Genetic algorithms simulate genetic systems and were first proposed by Fraser [8, 9] and Bremermann [2]. Despite the age of this early work, it still has relevance today, as it has recently been revisited by Fogel [6, 7]. John Holland, together with his students and colleagues at the University of Michigan in the 1960s and 1970s, is also credited with carrying out much of the pioneering work in GAs. Holland’s book of 1975 [11] is recognized as one of the seminal works in the area.

Genetic algorithms are computer programs that are based on Darwin’s principle of natural evolution. A GA holds a population of solutions (often referred to as chromosomes). The way in which each solution is represented is largely down to the designer of the GA. Historically, bit strings have been used but these are by no means the only possible representation. This work, for example, uses a set of real numbers that in turn represent the weights of a neural network. Each chromosome (shown in Figure 3) is assigned a fitness value that indicates the quality of the solution the chromosome represents.

Depth		Dominance		Discrimination weights						Sector weights									
Lateral	Vision	Lateral	Vision	Lateral			Vision			Lateral			Vision						
				Food	Prey	Pred	Food	Prey	Pred	0>0	0>0	1>0	1>1	0>0	0>1	1>1	1>1	2>0	2>1

Figure 3. An example of the chromosome structure used in E-BOIDS.

During the execution of a GA algorithm the population is continually replaced by new populations. The new populations are created by applying operators (crossover and mutation) to members of the existing population. An individual's chance of being chosen as a parent is proportional to its fitness. This is done so that the principle of natural selection is mimicked; that is, the fittest members of the population are allowed more opportunity to breed in the hope that they will pass their good genetic material to the next population. If this happens enough, the population should gradually improve as fitter and fitter individuals are created. Introductions to GAs can be found in Goldberg [10], and Michalewicz [15, 16].

### 3.3 Ecology

Evolutionary perspectives on schooling suggest that the environment plays a key role in its emergence, yet the number of factors to consider is potentially infinite. For example, the fluid dynamics as the E-BOIDS swim and displace the water could be considered. However, we have implemented only the principal components of the ecology.

From its conception, an organism will be either predator or prey and will recognize other organisms as such. Using predators and prey within the same system and evolving them together is important because one of the key reasons for schooling is the protection it provides.

At the start of its life, an organism has a numerical amount of energy, which it can spend on a number of behaviors. It is the implicit goal of every creature to pass on its chromosome to the next generation. Each type of organism is capable of the same repertoire of actions (Turn, Move, Eat, Collide, and Mate; note that mating is done asynchronously in the ecology) that use a predefined amount of energy (refer to Table 1). Since the amount of energy is used to decide which members are selected, it is advantageous to minimize energy expenditure. Predators eat prey and prey eat randomly placed food. If energy should fall below zero, a creature will die and turn into food for prey organisms.

The spatial distribution of food and organisms will be random for each generation. This is to ensure the population does not evolve to suit only one configuration of the world. Food will be placed in random clusters since PolyWorld [20] showed that territories of food could be very important to the ecology as they provide a common goal.

Table 1. Example of energy changes per behavior used in a typical simulation. Note that prey are not necessarily killed upon being (partly) eaten.

Behavior	Prey	Predator
Birth	+200	+200
Clash	-3	-3
Eat food	+20	+25
Eaten	-5	N/A
Mate	-2	-20
Move	-1	-1

There are a limited number of creatures within the world for each generation; hence, a selection scheme is necessary. The energy values are used as the fitness score: pairs of organisms, assuming that they had mated together during their lives, are selected as a function of the average energy value of the pair at the time of mating. The following is an outline algorithm for the evolution of generations:

1. Initialize population of chromosomes by randomly producing legal prey and predators, in a set ratio of 2:1 (prey:predator). Create ANN based on the chromosomes.
2. Run the creatures through a specified number of moves, and keep track of the pairs of organisms that have mated in the ecology. At the end of the generation:
  - 2.1. Give priority to fitter parents.
  - 2.2. Create new chromosomes by mating the current population using the genetic algorithm.
  - 2.3. Delete previous generation.
  - 2.4. Insert new members into the population.
3. Repeat stage 2 until no further generations are required (user intervention).

This strategy allows one to search in parallel through a large number of possible organisms, without discriminating against useful behaviors. A parameter is assigned that defines how likely crossover and mutation typically occur. Generally probabilities for mutation are 0.008 and crossover 0.6 [2].

#### **4 Implementation**

The main program (written in C++) controls the population dynamics such as evolution, movement, and the ecology for each generation. The output depicting the movements can be viewed using a separate system. The benefit of this approach is that the main program is not slowed down by graphics and that no artificial delays need to be implemented. Each generation can also be compared easily.

The population dynamics are displayed via VRML source code generated by the main program, which is parsed by an external VRML browser. This gives complete control over the viewing angle, speed, blurring, zooming, and coloring and enables dynamic navigation in the environment, all at the touch of a button. VRML is a platform-independent language, accessible over the Internet via the World Wide Web. In addition, HTML code was generated to produce statistics for each generation. This design enabled us to make all the results accessible online [5].

#### **5 Experiments with E-BOIDS**

This section describes experiments with the E-BOIDS world to investigate the effects of evolution on its inhabitants. There is no objective, quantitative measure of flocking or schooling [23]. As a consequence, many collective behavior studies resort to informal descriptions of the behavior of their system. However, several authors [4, 19] have recommended the use of standard statistical techniques to evaluate the behavior of synthetic systems. In an attempt to describe objectively the results in the following analyses, we have used a Euclidean distance measure to approximate schooling behavior and then subjected these measures to statistical methods [12]. In addition, a selection of behavioral results may be viewed online [5].



The first experiment looks at the ecology dynamics, that is, relationships developed in the environment between organisms. Highlighting these relationships, the next experiments aim to pick out common prey (Experiment 2) and predatory (Experiment 3) behaviors and the effects of changing parameters such as sensory configuration on behavior.

## 5.1 Ecology Dynamics

### 5.1.1 Method

A Euclidean distance measure was taken of each member of the population with respect to all the other creatures in the environment. The aim is to look at the development of global behavior in the system with each generation. Three variables were under consideration: first, the distance from food to each prey; second, the distance from each prey to another; and third, the distance from predators to prey.

Samples of 20 prey to 10 predators were taken, with 16 clumps of food (2 food items per clump). Energy values were set to encourage prey objects to be drawn toward food and predators to eat the prey. There were negative effects for collisions and for being eaten (Clash =  $-3$ , Eat food = 20, Eat prey = 25, Eaten =  $-15$ , Mate =  $-5$ , Move =  $-1$ ). Prey were set to swim slower than predators. The system was run for 25 generations, with 200 moves in each, and a total of 10 runs were made, therefore 1,500,000 moves were observed. The average distance for each generation was taken.

### 5.1.2 Results and Discussion

Figure 4 shows the average distance with each generation. It can be seen that with each generation predators learn to get closer to prey and that prey get closer together. Prey also seem to learn the value of food. To investigate the relations between these variables, a Pearson product moment correlation was performed.

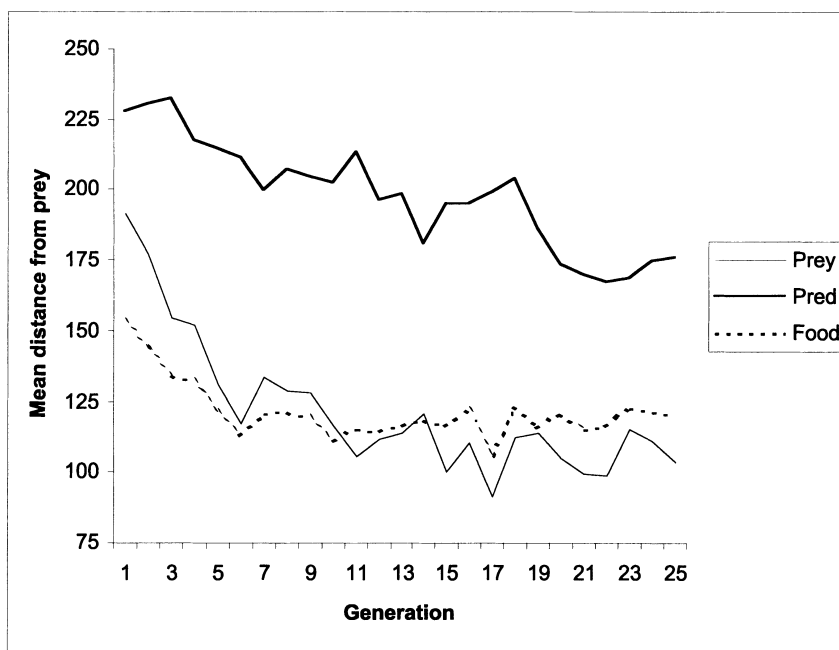


Figure 4. Distribution of distance with respect to prey for prey, predator, and food as a function of generation number.

Table 2. Means and standard deviations of distance from prey as a function of ecology type (with or without predators) and organisms (prey or food).

	Baseline distance			Average distance		
	Prey	Predator	Food	Prey	Predator	Food
Mean	152.29	184.44	174.01	121.71	197.68	121.39
Standard deviation	1.52	3.08	1.95	24.35	19.50	10.12

First, it was hypothesized that prey should adapt to schooling due to the influence of predators in the environment. The correlation between prey distance from prey and predators, respectively, supports this hypothesis (0.73,  $p < 0.01$ ). Second, it was hypothesized that, as prey get closer to other prey with each generation, they should get closer to food as well. The observed correlation (0.91,  $p < 0.01$ ) is consistent with this hypothesis. Lastly, it was hypothesized that predators will be drawn toward food since prey get closer to food. The observed correlation is 0.53 ( $p < 0.01$ ).

Another way to analyze schooling is to look at the average distance between organisms and prey and to compare it with the baseline distance, that is, the random distances between organisms and prey at the start of each generation. Measuring the change of these figures with time gives an indication of the behavior trends of the creatures. From Table 2, it appears that prey tend to get closer together, that predators are not as close to prey on average than they are at the start of the simulation, and that prey are closing in on food objects. The standard deviation indicates quite a variation in the development. To clarify the situation, some statistics are calculated to check whether the behaviors observed are significantly different from chance.

The first analysis looks at the prey dynamics. We can describe our experiment as consisting of two groups (“ecologies”), one where the distance is measured at the start of the simulation, and one where the average distance is taken. For each of these groups, measurements are repeated over 25 trials (generations). A  $2 \times 25$  mixed analysis of variance (ANOVA) was therefore conducted, with ecology type as a between-group variable with 2 levels (baseline and average), and generations as a within-group variable with 25 levels. The dependent variable was the distance between prey E-BOIDS.

There was a significant main effect of ecology type ( $F_{1,18} = 9.39$ ,  $MS_e = 116,953.69$ ,  $p < 0.01$ ), a significant main effect of generations ( $F_{24,432} = 5.45$ ,  $MS_e = 2,981.72$ ,  $p < 0.001$ ), and a significant interaction between ecology and generations ( $F_{24,432} = 5.43$ ,  $MS_e = 2969.39$ ,  $p < 0.001$ ). Overall, it appears that the spatial distribution of prey organisms is different from their random starting positions, and that with each generation, there is a significant change. Figure 5 shows the distribution of these results (left panel).

The next analysis aims to look at the relationship between predators and their prey. A  $2 \times 25$  mixed ANOVA was conducted with ecology type as a between-group variable with 2 levels (baseline and average) and generations variable with 25 levels. The dependent variable was the distance between prey and predator E-BOIDS.

There was no significant main effect of ecology type ( $F_{1,18} = 2.84$ ,  $MS_e = 21,923.77$ ,  $p > 0.05$ ), a significant main effect of generations ( $F_{24,432} = 2.93$ ,  $MS_e = 2,173.98$ ,  $p < 0.001$ ), and a significant interaction between ecology and generations ( $F_{24,432} = 2.32$ ,  $MS_e = 1,722.77$ ,  $p < 0.01$ ). These results suggest that, on average, the distance between prey and predators after several generations is not different from the baseline distance. However, the significant interaction as well as the significant main effect of generation show that predators tend to reduce their distance to prey, as shown in Figure 5 (right panel). Comparisons were not performed on the distance between prey and food because the food is randomly repositioned in each generation and after 50 moves.

It can be concluded that prey and predators are developing to show behavioral traits. Prey are moving closer to food and appear to be inclined to go toward each

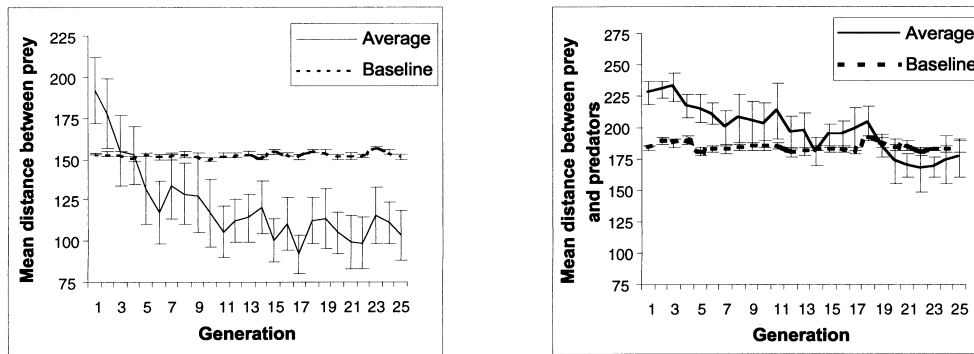


Figure 5. Behavior of prey E-BOIDS in different ecologies. Left graph shows the mean distance between prey. Right graph shows the mean distance between prey and predators. Error bars indicate the standard errors of the mean.

other. Predators appear to be going toward prey objects, though not significantly so from their baseline distance. However, it must be remembered that the baseline is only a snapshot of a random starting position, not random movement.

The results have demonstrated that a simple measure such as the Euclidean distance can be used to quantify schooling behavior. Some important relationships in the ecology have been identified that may not have been obvious from the outset. They indicate that prey increasingly move together and that predators appear to be drawn to food areas. In fact there are a range of behavioral repertoires observed. It appears that there is a sensory overload around food areas. This means that predation is more successful outside food areas, and that it is best to steer away from them. But at the same time some predators are drawn toward food in the knowledge that prey are likely to be around. A balance between these behaviors is generally observed, some predators driving the prey out of the food areas while the others stay back to hunt them down.

## 5.2 Prey Dynamics

The results of the ecology dynamics (Experiment 1) look promising, but they were obtained using a rather primitive measure of schooling behavior. To understand fully whether schooling really occurs, further tests must be performed. This section focuses on the dynamics of prey behavior.

### 5.2.1 Method

Mataric [14] states that flocking is the combination of a number of basic primitives: collision avoidance, following, dispersal, aggregation, and homing. The aim of this experiment is to look for evidence of each of these behaviors. To perform this experiment, we attempted to increase all the energy constants associated with moving close together. Therefore mating caused an increase in energy, eating vegetation had no effect, and a clash had a small effect. The system was run for 25 generations with 200 moves in each, and a total of 10 runs were made. The average distance for each generation was taken. As expected, schooling behavior was observed.

### 5.2.2 Results and Discussion

The experiment was run over many generations, with varying parameters to look for evidence of schooling behavior. Unfortunately, observation was the only measure available. Selections of the behavioral results are shown in Plate 1.

The results appear to satisfy Mataric's primitives and demonstrate schooling behavior. Each prey appears to value the school more than food itself, with the security that it



Plate 1. A selection of the schooling behavior observed in the prey E-BOIDS. Blue indicates prey, red predators, and green food.

provides. The motion of the school is due to slight variation in the distance senses of each creature, each guiding the other. Some creatures appear to be providing a “global lookout” whereas the others maintain their close-up motion, guiding them toward food objects and toward each other. The lateral line is always used as a repulsion force keeping each other apart to avoid collisions. Figure 6 shows a sensory network taken from a typical prey in the final generation.

The weights in the lateral line indicate that the prey will generally turn away whereas the weights in vision generally guide them toward their target. Prey tend to be guided toward food (0.70) in the vision system, whereas in the lateral range, greater attention is given to predators (0.43).

In contrast, a successful predator network is shown in Figure 7. The lateral line and vision of the predator serve a completely different function to the prey’s. Here predators employ the sensory system to guide them toward prey and food in the environment. Again we see a high weighting of the lateral sense (0.51) but this time a small depth, perhaps due to the high-resolution fields required to pick out their prey. It is interesting to note that predators regard each other more highly than their prey, perhaps due to the fewer numbers of predator to prey.

To test the differences in weights between prey and predators, we used a matched *t*-test. We found reliable differences ( $p < 0.05$ ) for dominance (both lateral and vision) and for depth of vision. With respect to attention, we found reliable differences for food and predators in the lateral line, and in all the sensory weights of the lateral line. The role of the lateral line would therefore appear to be one of the defining characteristics in the difference between predators and prey.

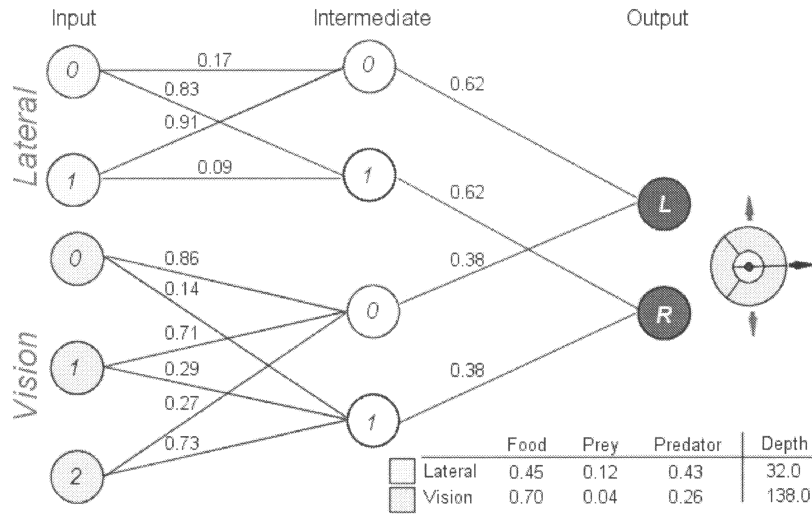


Figure 6. Evolved sensory network of prey. The figure indicates the number of input units for each sense, and how they map to the output nodes (turn left, turn right). The network feeds from left to right; input units show the lateral and visual sensors. The weights of food, prey, predators and the depth of the senses are shown in the table (bottom right). Intermediate nodes modulate the balance between lateral and vision sensor input (dominance). The structure of an E-BOID is shown right.

The number of visual and lateral sectors were varied, but there appears to be very little difference in the number required to school. The prey dynamic results were collected using two lateral and three visual sectors. It would appear therefore that prey organisms need only very low-resolution visual systems to be successful in their environment. The issue is taken up later when predator visual systems are discussed.

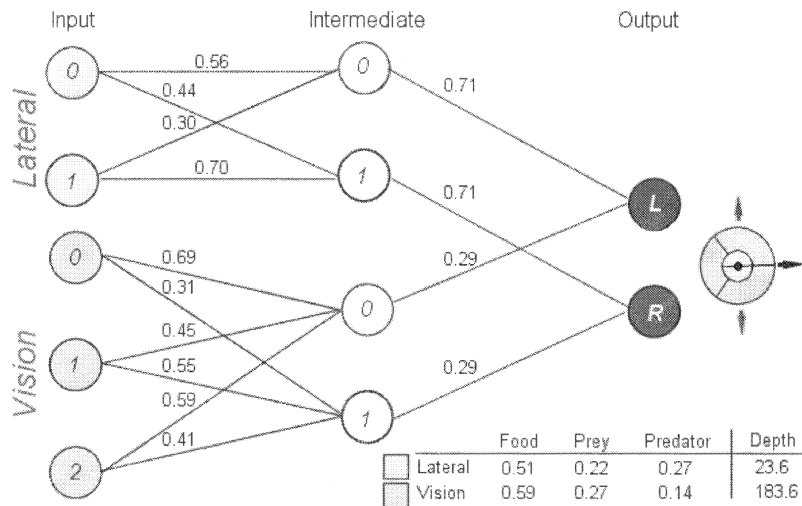


Figure 7. Evolved sensory network of predator.

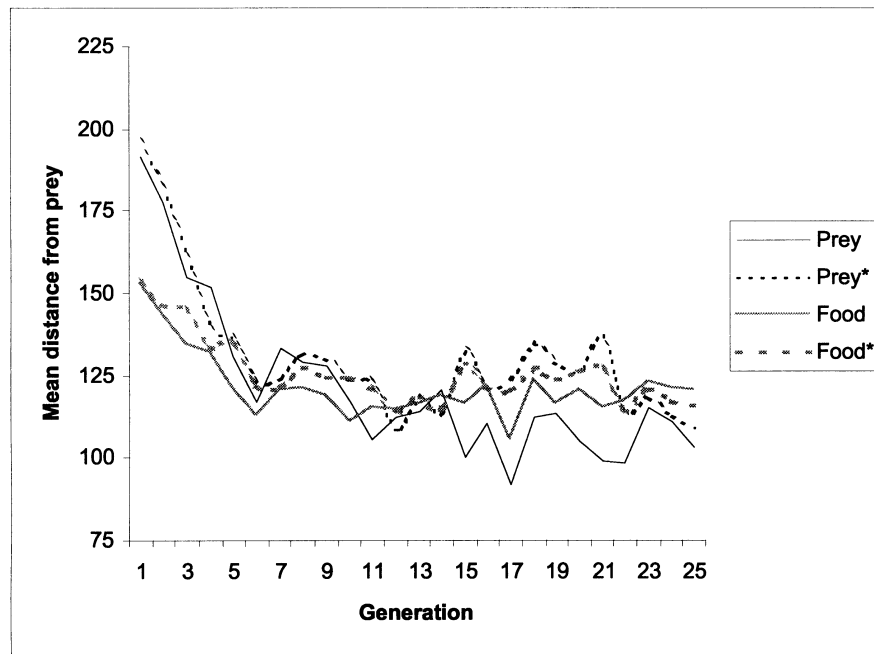


Figure 8. Distribution of distance over each generation for runs with and without (dotted) predators.

### 5.3 Predator Dynamics

#### 5.3.1 Method

The initial results from our ecosystem suggest that predators were encouraging prey to school. This section explores whether this was the case. The first experiment is repeated with and without predators, to see if predators are a necessary requirement of the ecosystem; conditions are kept the same (Clash =  $-3$ , Eat food =  $20$ , Eat prey =  $25$ , Eaten =  $-15$ , Mate =  $-5$ , Move =  $-1$ ). The system was run for 25 generations, with 200 moves in each. A total of 10 runs were made. The average distance for each generation was taken.

#### 5.3.2 Results and Discussion

Referring to Figure 8, it would appear that the pattern of the results is quite similar, though it can be seen that the simulation “with predators” shows more schooling behavior, as operationalized by the distance from other prey. This result demonstrates the role of schooling in goal-directed behavior such as eating, with each member of the school aiding the others, and the role of predators in shaping the school.

As in the first experiment, we also analyze the results using average distance across generations. The results (Table 3) suggest that predators have a part to play in the environment, but perhaps not as important as might have been thought. To establish

Table 3. Means and standard deviations of distance from prey as a function of ecology type (with or without predators) and organisms (prey or food).

	With predators		Without predators	
	Prey	Food	Prey	Food
Mean	121.71	121.39	130.89	125.56
Standard deviation	24.12	19.20	21.03	10.18

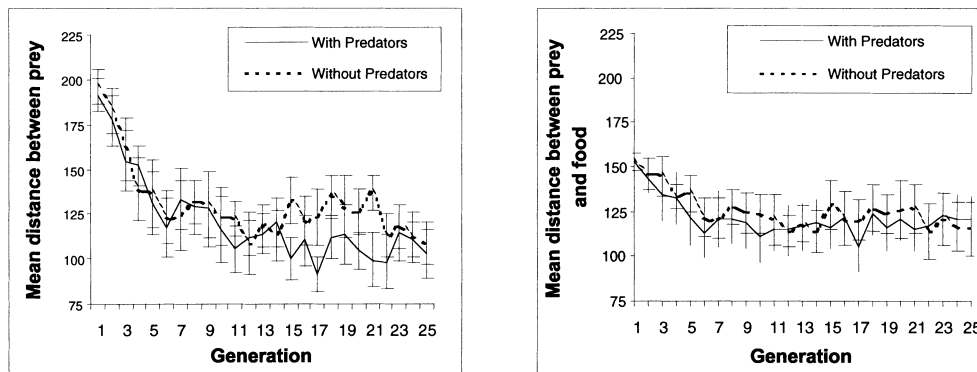


Figure 9. Distribution of distance between prey with respect to generations, under two ecology types (with and without predators). Error bars indicate the standard errors of the mean.

whether the effect is reliable, a  $2 \times 25$  mixed ANOVA was conducted with ecology type as a between-group variable with 2 levels (with and without predators) and generations as a repeated measures variable with 25 levels. The dependent variable was the distance between prey E-BOIDS.

There was no significant main effect of ecology type ( $F_{1,18} = 0.47$ ,  $MS_e = 10,323.96$ ,  $p > 0.05$ ), a significant main effect of generations ( $F_{24,432} = 7.83$ ,  $MS_e = 9,573.73$ ,  $p < 0.001$ ), and no significant interaction between ecology and generations ( $F_{24,432} = 0.65$ ,  $MS_e = 798.87$ ,  $p > 0.05$ ). Overall, evolution appears to have an effect on the distance between prey, yet the ecology type does not have a significant role. Figure 9 (left panel) illustrates this result.

Next, the relationship between prey and food is considered. A  $2 \times 25$  mixed ANOVA was conducted with ecology type as a between-group variable with 2 levels (with and without predators) and generations as a repeated measures variable with 25 levels. The dependent variable was the distance between prey E-BOIDS and food.

There was no significant main effect of ecology type ( $F_{1,18} = 0.08$ ,  $MS_e = 2,142.40$ ,  $p > 0.05$ ), a significant main effect of generations ( $F_{24,432} = 3.44$ ,  $MS_e = 1,877.25$ ,  $p < 0.001$ ), and no significant interaction between ecology and generations ( $F_{24,432} = 0.34$ ,  $MS_e = 185.38$ ,  $p > 0.05$ ). Overall it appears that evolution, again, has a strong effect, this time on the distance between prey and food, and that, again, ecology type does not have a significant effect. Figure 9 (right panel) shows the distribution of these results.

To analyze the change of behavior with each generation further, a regression analysis was performed on different ecology types (with and without predators) to test whether one generation learns to school more quickly than the other. It was hypothesized that aggregation follows a power law. Figure 10 shows the distribution of the data and the fitted power function. It was found that the distance between prey does indeed follow a power law; both with predators ( $F(2, 23) = 124.88$ ,  $p < 0.001$ ) and without predators ( $F(2, 23) = 52.69$ ,  $p < 0.001$ ). Interestingly when predators are in the system there is a steeper power law, as reflected by the differences in exponent ( $-0.1995$  vs.  $-0.1445$ ).

The results suggest that predators have a role in the emergence of schooling behavior in the ecosystem; however, this role is perhaps not as strong as might have been hoped. Even so, we believe that this is purely a reflection of the measure used rather than behavior. To illustrate the point, Plate 2 shows the predation observed.

It appears that predators experience a sensory overload when faced with a school, yet when the prey E-BOIDS split from the group the predators can home into them. To illustrate the difference in the sensory systems seen when predators are not in the system, a typical network of a prey is shown in Figure 11.

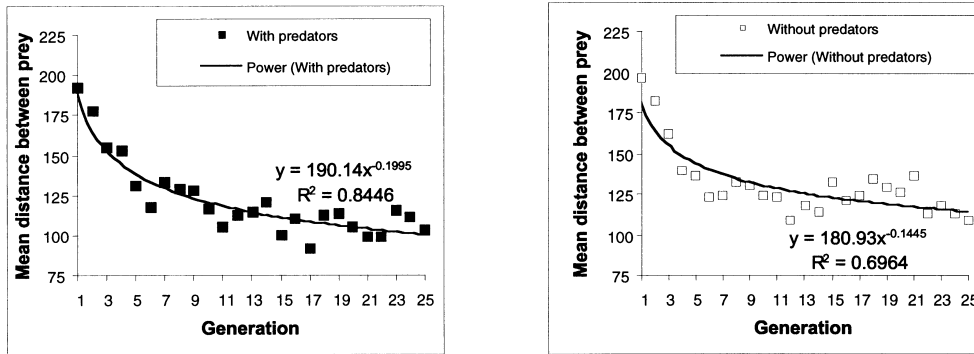


Figure 10. Behavior of prey E-BOIDS in different ecologies. Left graph shows the mean distance between prey when predators are present in the environment. Right graph shows the mean distance between prey in the absence of predators. The curve shows the best-fitting power function.

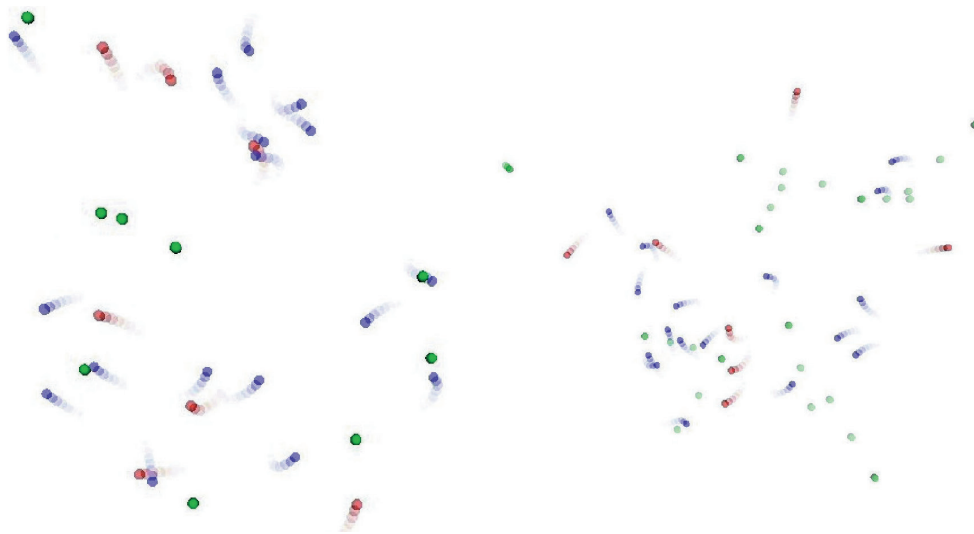


Plate 2. Evolution of predatory behavior in the system. Blue indicates prey, red predators, and green food. Of particular note is that prey organisms are “swimming” away from predators and predators are tracking prey.

In contrast to Experiment 2, Figure 6, the creature is highly oriented toward everything with high weights on food in vision (0.55) and very little on predator. The lateral sense dominates over vision, in terms of attention (0.85 vs. 0.15) and depth (79.4 vs. 138). This indicates that the prey are not looking out for predators as they were before; now, they consume as much food as possible to increase their energy to compete for the next generation. As with the experiment outlined in Section 5.2 we used a matched *t*-test to identify reliable differences in the weights of the networks of prey with different ecologies. We found statistically reliable differences ( $p < 0.05$ ) for all weights except for the attention weight prey (vision) and lateral weights ( $1 \rightarrow 0$ ,



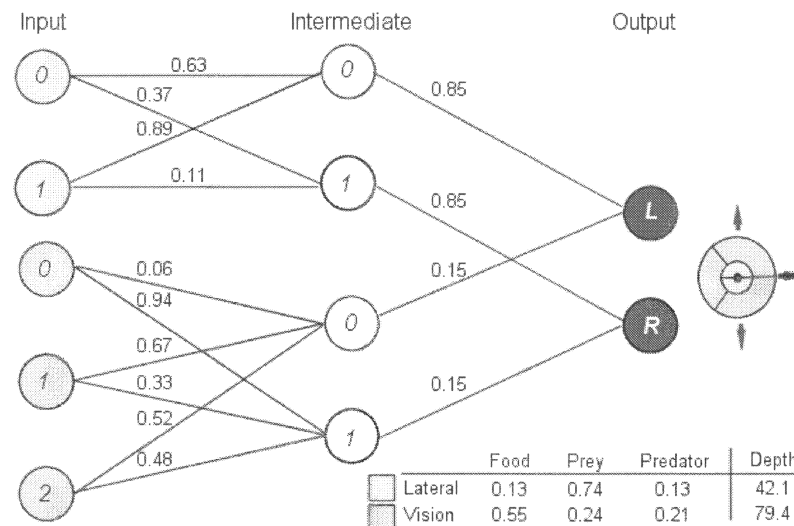


Figure 11. Sensory network of an evolved prey from an ecology with no predators. The figure indicates the number of input units for each sense, and how they map to the output nodes (turn left, turn right). The network feeds from left to right; input units show the lateral and visual sensors. The weights of food, prey, and predators and the depth of the senses are shown in the table (bottom right). The structure of the E-BOID is shown right. Contrast this network to that found with predators (see Figure 6).

1  $\rightarrow$  1; refer to Figure 11 for annotations), and visual (2  $\rightarrow$  0, 2  $\rightarrow$  1). There were marginally significant differences for the lateral food weight ( $p < 0.06$ ) and the visual food weight ( $p < 0.08$ ).

To explore schooling behavior further, the number of visual and lateral sectors were varied; it was found that increasing the number of visual sectors produced superior predatory behavior. For example, it appears that frontal visual systems in predators are very important in directing their attention, whereas prey require a peripheral visual system to detect the objects all around them. To aid predators, three visual sectors were used, with each covering  $120^\circ$ : the majority of sectors dominated frontal vision because  $240^\circ$  were focused forward to the left and right with  $120^\circ$  looking backward. This observation is consistent with evolutionary work that suggests predators have higher resolution visual systems than prey [3].

The results have also shown the importance of coevolving creatures rather than artificially introducing hand-crafted prey-targeting systems [22]. Designing a controller that could foresee, or indeed cope with these interactions would be a hugely complex task, were it to be hand coded.

## 6 Summary and Discussion

This article has considered the role of sensors in the emergence of primitive collective behaviors. We identified a particular species (the fish), looked at the role of ecology in the evolution of sensors, and designed a flexible structure that would enable complex movement and perception on a simple base network. With schooling in mind we hoped to extend previous work that had failed to produce these results.

Primitive behaviors were evolved through an ecology in which predator and prey coexisted within the same environment. Although these behaviors have been achieved with a fairly simple base network on which all creatures exist, the simplicity has not constrained the behavioral repertoires observed. Using the same base network, two

highly different behaviors were observed, that of predatory and prey behavior. These behaviors were not hand coded, in contrast to previous studies; rather, they were evolved using an ecology. Analyzing the role of the ecology showed that predators are of key importance as a means of encouraging prey to school. Schooling behavior was achieved, together with predation, and these results were validated with statistical analysis. These results have shown that small changes to sensory systems can produce major changes in behavior. In that respect, our work contrasts with previous studies, which exclusively relied on informal descriptions to evaluate the success of the simulations, without taking into account the effect of random variation. The work also improves on previous research, in that these behaviors have been produced with a simple dynamic neural network structure without a fitness function incorporating a measure of schooling [23].

One way to improve the system is to increase its complexity. For example, as a form of motivation, it may be useful to represent the energy level as an input. When an E-BOID would get low on energy, it would be encouraged to forage for food [23]. However, it is unclear how this energy level would feed into the behavioral system. One possibility is to use a Boltzmann machine controller that could be evolved to motivate it. Adding such new complexity may, however, violate one of the tenets of ALife, which is to capture the essence of a system.

Many researchers appear to design complexity into their representations and, in doing so, constrain the generalizations that might be made. By contrast, the work described in this article highlights the idea that complex behavior can be achieved by simple interactions, and that many of these subtle developments cannot be foreseen using conventional techniques.

Evolutionary methods are a powerful programming technique. It would have been difficult, if possible at all, to foresee or even implement the results reported in this article using hand-crafted production rules. Artificial neural networks created a vast array of behaviors, including schooling, which has proved very problematic to achieve in related work. Yet the ANN used in E-BOIDS was not as complex as many of these unsuccessful systems. The reason for success is due to the simple base network and to the artificial ecology, which together overcame the difficulty in assigning fitness to the population to evolve schooling behavior.

### Acknowledgments

We thank Peter Lane, as well as two anonymous reviewers, for their detailed comments on this paper.

### References

1. Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
2. Bremermann, H. J. (1958). *The evolution of intelligence. The nervous system as a model of its environment*. (Tech. Rep. No. 1, Contract No. 477(17)). Seattle: Department of Mathematics, University of Washington.
3. Cliff, D., & Miller, G. F. (1996). Co-evolution of pursuit and evasion II: Simulation methods and results. In P. Maes, M. Mataric, J. A. Meyer, J. Pollack, and S. W. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 506–515). Cambridge, MA: MIT Press.
4. Cohen, P. R. (1995). *Empirical methods for artificial intelligence*. Cambridge, MA: MIT Press.
5. E-BOIDS web pages. (2000). <http://www.cs.nott.ac.uk/~gwk/alife/ward2001>.
6. Fogel, D. B., & Anderson, R. W. (2000). Revisiting Bremermann's genetic algorithm: I. Simultaneous mutation of all parameters. In A. Zalzal et al. (Eds.), *Proceedings of Congress*

- on *Evolutionary Computation 2000 (CEC 2000)* (pp. 1204–1209). Piscataway, NJ: IEEE Service Center.
7. Fogel, D. B., & Fraser, A. S. (2000). Running races with Fraser's recombination. In *Proceedings of Congress on Evolutionary Computation 2000 (CEC 2000)*. La Jolla, CA, pp. 1217–1222.
  8. Fraser, A. S. (1957). Simulation of genetic systems by automatic digital computers. II. Effects of linkage on rates under selection. *Australian Journal of Biological Science*, *10*, 492–499.
  9. Fraser, A. S. (1960). Simulation of genetic systems by automatic digital computers. IV. Epistasis. *Australian Journal of Biological Science*, *13*, 329–346.
  10. Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.
  11. Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor: University of Michigan Press.
  12. Howell, C. H. (1987). *Statistical methods for psychology*. Boston: PWS-KENT.
  13. Husbands, P., Harvey I., Cliff, D., & Miller, G. (1997). Artificial evolution: A new path for artificial intelligence? *Brain & Cognition*, *34*, 130–159.
  14. Mataric, M. J. (1992). Designing emergent behaviors: From local interactions to collective intelligence. In J. A. Meyer, H. L. Roitblat, and S. W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 432–441). Cambridge, MA: MIT Press.
  15. Michalewicz, Z. (1996). *Genetic algorithms + Data structures = Evolution programs* (3rd ed.). New York: Springer.
  16. Michalewicz, Z., & Fogel, D. B. (2000). *How to solve it*. New York: Springer.
  17. Partridge, B. L., & Pitcher, T. J. (1979). Evidence against a hydrodynamic function of fish schools. *Nature*, *279*, 418–419.
  18. Partridge, B. L. (1982). The structure and function of fish schools. *Scientific American*, *246*, 90–99.
  19. Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. Cambridge, MA: MIT Press.
  20. PolyWorld. (1999). <http://alife.santafe.edu/alife/software/polyworld.html>.
  21. Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. [SIGGRAPH 1987 Conference Proceedings] *ACM Computer Graphics*, *21*, 25–34.
  22. Reynolds, C. W. (1992). An evolved, vision-based behavioral model of coordinated group motion. In J. A. Meyer, H. L. Roitblat, and S. W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 384–392). Cambridge, MA: MIT Press.
  23. Werner, G. M., & Dyer, M. G. (1992). Evolution of herding behavior in artificial animals. In J. A. Meyer, H. L. Roitblat, and S. W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 393–399). Cambridge, MA: MIT Press.
  24. Zaera, N., Cliff, C., & Bruten, J. (1996). (Not) evolving collective behaviors in synthetic fish. In P. Maes, M. Mataric, J. A. Meyer, J. Pollack, and S. W. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 635–644), Cambridge, MA: MIT Press.