Predation as a Motivation for Schooling Behavior with Simulated Fish

Melissa O'Connor and Joe Boninger
Department of Computer Science, Swarthmore College

May 9, 2014

Abstract

It has been theorized in biology that one of the primary motivations for schooling behavior in fish is escaping predators. When potential prey form larger schools, the predator can become confused and have difficulty trying to target one individual prey from the larger mass. We attempt to model this phenomenon by designing a custom graphical simulator with two different types of simulated robots—prey and predator. The predator is hard-coded to pursue prey that are not in schools and to pursue a random prey at each time step if all prey are in schools. The simulated prey are controlled by a NEAT artificial neural network with fitness based on how many fish are able to avoid being "eaten" by the predator and how long the school as a whole is able to survive. NEAT (NeuroEvolution of Augmenting Topologies) is an evolutionary robotics tool for evolving artificial neural networks (ANNs). NEAT begins with an ANN without any hidden nodes and, as the experiment proceeds, adds hidden nodes as needed. We found that by using a neural network, the prey are able to learn to school successfully within ten generations, even with a small population. These findings suggest that a predator can be good motivation for prey to school and that the schooling task is easier to solve than we originally thought.

1 Introduction

1.1 Background

Many different motivations have been proposed to explain schooling behavior in biological systems such as fish or birds. One of the most popular and common theories is that schooling is a way to confuse predators. According to *How It Works*, an effect known as the predator confusion hypothesis is responsible for this behavior:

> By grouping into a tight, regimented pattern, the fish minimise their chance of being picked off by generating a sensory overload to a predator's visual channel. The swirling mass of twisting silvery fish creates a blending effect where the predator struggles to track a single target and becomes confused [1].

We chose to model schooling in the presence of a predator by evolving neural network controllers for simulated robots through an evolutionary robotics tool called NeuroEvolution of Augmenting Topologies, or NEAT [4]. NEAT is a specialized type of artificial neural network, or ANN. An ANN is a data structure made up of nodes and connections that is loosely inspired by the structure of neurons in a brain. One layer of nodes is the input layer, which takes in input. Another layer of nodes, the output layer, outputs data depending on the input, the structure of the nodes, and the weights of the connections between them. There also may be one or more hidden layers of nodes between the input and output layers that transform the data. Over the course of evolving a neural network, the weights of connections are changed until any given input produces the desired output.

NEAT begins with a simple topology consisting of only an input and output layer, but it can add hidden nodes over the course of evolution if the complexity of the problem demands it. Nodes and connections both change over the course of evolution. NEAT is able to preserve structure when mating networks by tracking the placement and order of creation of each node using historical markers. The markers help align analogous nodes from different brains. We apply NEAT in a custom-built Python graphics simulator which will be further detailed later in this paper.

1.2 Related Work

There have been many experiments performed that deal with schooling behavior in a prey and predator situation. In *Predator confusion is sufficient to evolve swarming behavior*, Randal Olson, et. al. model a two-dimensional continuous environment in which there is coevolution of the predator and prey behaviors [2]. Like in our experiment, prey are able to move forward or rotate on each time step. Unlike our experiment, however, their prey also have the option to stay still. In their experiment, the prey evolve several distinct behaviors, including a "bait ball"-like behavior which we observed in our experiments. A "bait ball" occurs in biology when fish all swim in a tightly packed spherical formation, often as a result of being threatened by predators. The fitness functions used in this experiment are based solely on either the number of prey eaten or the number of prey that survive; in our fitness function, we added an additional metric that rewards prey based on the number of time steps survived.

Another experiment that evaluates the role of the predator in prey schooling behavior is described by Ward, Gobet, and Kendall [5]. Like Olson, these authors perform coevolution of the predator and prey. They incorporate an additional element, a food source, to more closely model an ecosystem. The results show that, even with a food source, the predator is still instrumental in encouraging schooling behavior among the prey. There were two types of predator behaviors observed in this experiment: one where the predators surrounded the food source and another in which they directly attacked the prey.

As inputs to the neural networks, Ward, Gobet, and Kendall incorporate values that represent the distance and angle to the closest neighbors, just as we do with the three closest neighbors. This experiment uses a standard genetic algorithm, whereas we use NEAT for evolution. We chose not to add a food source because we did not want the schooling behavior to be a result of all agents travelling towards the same location. In our experiment, the schooling behavior emerged solely due to the presence of the predator and the threat of being attacked and "eaten".

Pugliese and Marocco also perform experiments that involve flocking [3]. They focus on the emergence of leadership patterns as well as what happens when a predator is introduced into an environment. Additionally, they experiment with allowing the robots to communicate via sound. The predators are motivated to hunt for prey while the prey are motivated to escape from predators. A neural network was evolved using evolutionary robotics feed-forward methodology with a fixed topology of nine input nodes, two hidden nodes, and two output nodes.

Pugliese and Marocco obtain an interesting result: instead of their prey flocking, their predators tended to flock. The prey do not group in any particular manner; they just explore the environment and avoid the predators when they are nearby. We only have one predator present in our simulation and observe behavior in which the prey flock rather than the predator. Our result could be due to our predator moving faster with respect to the prey than Pugliese's predators; our prey are not able to outrun the predator. Pugliese and Marocco also observe correlations between aggregation and fitness, as well as leadership and fitness. These indicate that the more aggregated groups with more leadership were also more likely to obtain higher fitness. Leadership behavior is observed when there is a clear relationship between a leader of the group and the rest of the flock; this is the opposite of peer-to-peer behavior.

2 Experiments

Our experiments were run using a custom-built 1,000 x 1,000 pixel Python simulator. We factored twenty percent noise into the prey's neural network inputs to better model reality. In biology, it is unlikely that prey would have perfect data about their surroundings. Robot sensors are also noisy, which makes these results more applicable to physical robots. Additionally, the simulation environment is bounded — when a prey or predator approaches the edge of the bounding box, it bounces in the opposite direction in which it was headed.

We ran twenty generations, each with a population size of fifteen neural networks. One generation consisted of three trials in which there were 15,000 time steps; the predator was immobile for the first 3,000 time steps, allowing the prey some time to school. It then started moving after this grace period.

Each neural network controlled a school of twenty prey. The prey were placed at random locations and with random headings within the simulation environment. Each trial, the predator started with a northwest heading in the bottom-righthand corner of the window. Upon being "eaten" by the predator, the prey are removed from the simulation. As mentioned before, the prey are unable to outrun the predator; their maximum speed is less than the predator's maximum speed. The prey must also maintain a speed above zero.

We defined a school as a fish having at least two neighbors present within a specified radius of itself. Note that this implies that a fish could be in a school while one of its neighbors may not be. This also means schools must consist of at least three fish.

Figure 1 depicts the environment. The predator, or "shark," is in the lower righthand corner; the black circle indicates his current heading. The prey, or "fish," are the smaller circles and are colored depending on whether or not they are properly schooled. The orange fish are not in a school, while the purple fish are in a school.

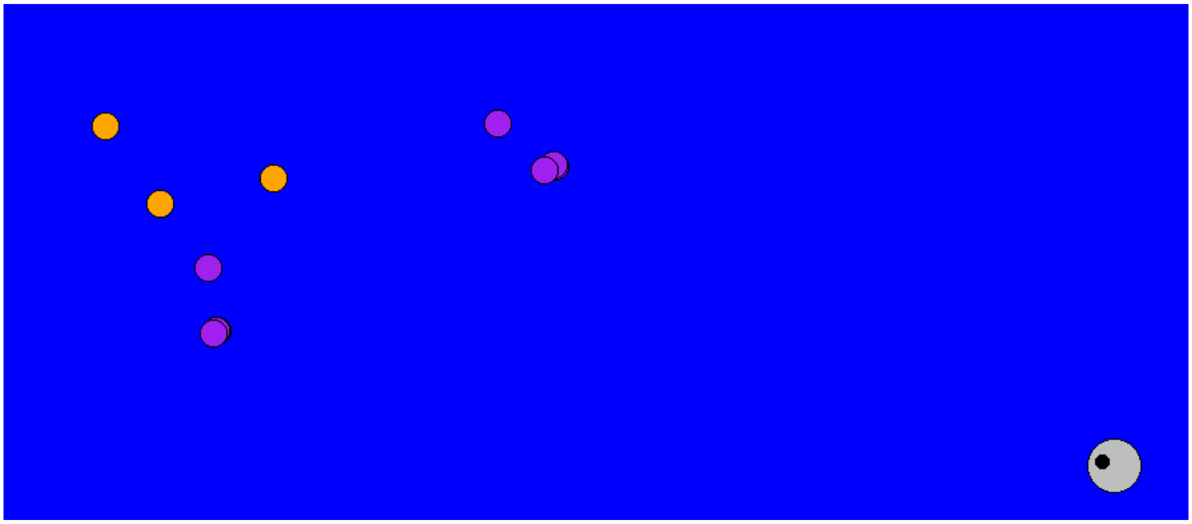

Figure 1: The Python simulation environment. The predator, or shark, is the large grey circle, whereas the prey, or fish, are the smaller circles. A fish is purple if it is in a school and orange if it is not in a school.

2.1 Neural Network

Neural network controllers were evolved with NEAT selecting for individuals that are better at surviving in the presence of a predator. The starting network topology, as seen in Figure 2, contained ten input nodes and two output nodes.

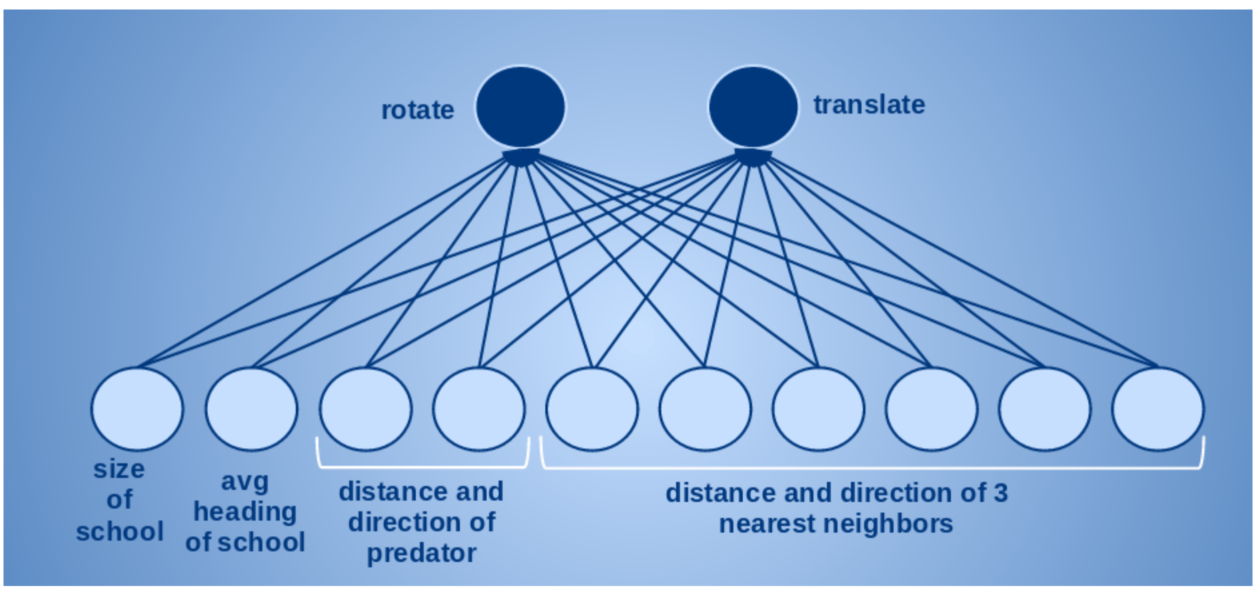

Figure 2: Starting network topology for prey, with labelled inputs and outputs.

The inputs to the neural network are intended to give the prey information about the predator as well as the prey's own school and whether or not it is schooling. The first two inputs give information about the school; they are the size of the school and the average heading of the school. The next two inputs provide the distance and direction of the predator. The last six inputs all give the distance and direction to the three nearest neighbors. All inputs were scaled between 0 and 1 and all inputs, besides the size of the school, have up to twenty percent noise incorporated into them.

Direction refers to the minimum angle difference between the prey's current heading and the heading the prey would need to have to face the nearest neighbor (or in the case of the predator, to face the predator). Angles scaled between 0 and .5 indicate a counter-clockwise turn, while angles scaled between .5 and 1 indicate a clockwise one. Angle differences were calculated by moving the coordinate plane of the window to put the origin on the prey, and then rotating the plane with a rotational matrix so that the heading of the prey lined up with the y-axis.

The two outputs of the neural network were rotate and translate values which determined how the prey moved at each time step.

Fitness was calculated as an average over three trials, and was a value scaled between 0 and 1. The fitness values were assigned according to the following fitness function:

$$0.4 * \kappa + 0.6 * \alpha, \tag{1}$$

where κ is the number of time steps the school survives after the predator begins moving and α is the number of fish that survive. Both κ and α were scaled between zero and one. The effect of this is that forty percent of the school's fitness is based on how long the school survives and sixty percent is based on the how many fish in the school survive.

2.2 Predator

The predator in our experiment was hard-coded and followed an algorithm that determines which prey it should "attacked." If the predator touches a prey not in a school, that prey is automatically eaten and removed from the simulation. If a predator touches a prey in a school, there is a one in five hundred chance of that prey getting eaten at each time step.

Pseudocode follows which describes the algorithm that determines the predator's behavior:

```
if there are prey not in schools:
     target = closest prey not in a school
else:
     target = a random prey from the set of all prey

if angle difference between heading and target is small:
     move forward
else:
     turn heading in direction of target
```

Note that in the second case, where all prey are in schools, the predator chooses a different target at each time step. It will only move toward this target if its heading is already pointing toward it—otherwise, it will rotate. This slows the predator down, but still allows it to catch prey by gradually moving in their direction.

Having the predator randomly target prey when they are in schools was implemented to better model biology. It is analogous to the predator's visual input being overloaded with sensory input from a school of prey.

3 Results

In two out of three runs of our experiment, the simulated prey were able to learn to school. Even with our small population size, convergence to a solution typically happened around the tenth generation. Figure 3 shows the best fitness and average fitness for a successful run. Some successful neural networks used as few as one hidden node, further suggesting that schooling and avoiding the predator is not necessarily a complicated task to learn. Figure 4 shows one example of a successful phenotype that contains one hidden node.

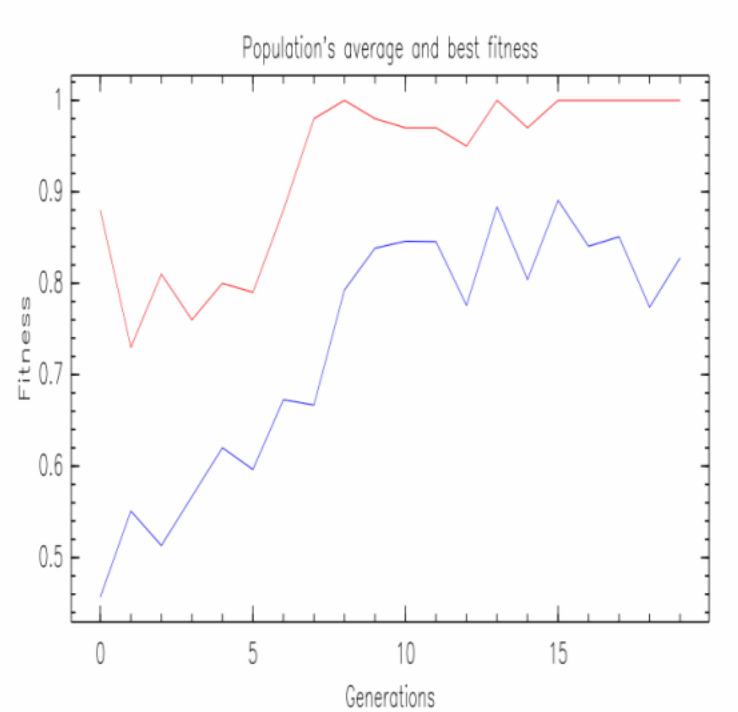

Figure 3: Graph of fitness results from the first run of our experiment. The upper line indicates the best individual from the population, while the lower line represents the average fitness.

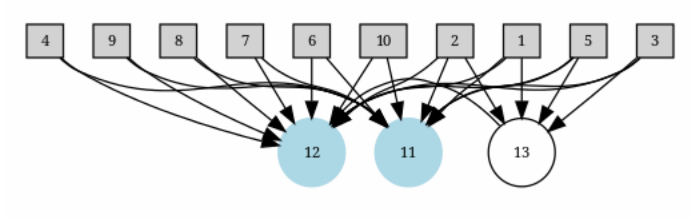

Figure 4: The resulting phenotype from generation 18 in the first run of our experiment. Nodes 1-10 are the input nodes, 11 and 12 are the output nodes, and 13 is the hidden node added during evolution.

In each successful case, the prey learned a similar method for achieving high fitness; in the beginning of each trial the prey form groups of three or more by moving toward their nearest neighbors. After they are in schools, they oscillate in circles without traveling. Oscillation occurs because we prevented them from remaining stationary between time steps.

When the predator comes close to a school of prey, the prey move together in what seems to be an arbitrary direction not toward the predator. Often the prey had trouble moving as a group away from the predator, frequently moving in patterns that circled back on themselves. These two features are interesting: first, because the most obvious direction for the prey to move would be directly away from the predator. However, they seem to move together in an arbitrary direction. How exactly they choose the direction or coordinate simultaneous movement is unclear—it could be the average heading of the school. Figure 5 shows the fish travelling away from the predator as it approaches their school. Second, the fact that the group movement away from the predator tends to double back on itself demonstrates that this was likely the most complex part of the schooling task. Even the successful schools had trouble moving together in a consistent direction. However, because the predator chooses a random target at each time step when pursuing schools of fish, it was unlikely to pursue one school for a sustained amount of time. Thus, we were unable to observe the schools of prey moving together, away from the predator, for long periods of time.

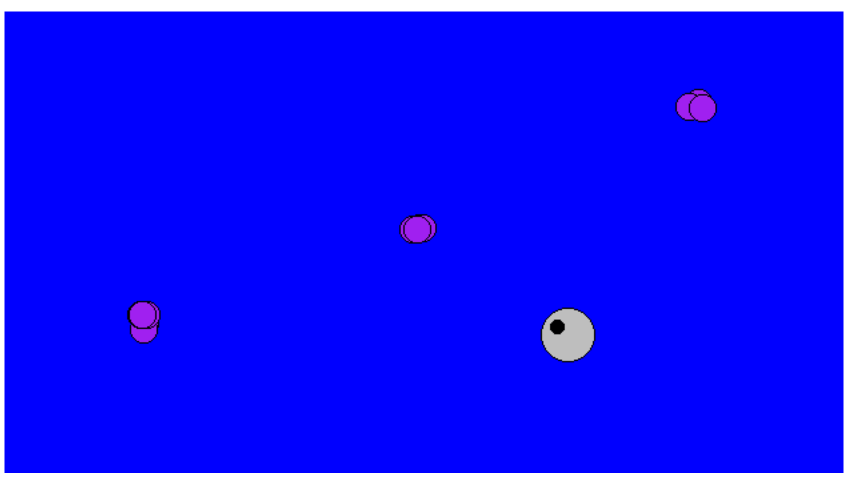

Figure 5: The schooled fish successfully escaping from the shark.

4 Discussion

Our experiment was a success, and our hypothesis was supported. By hard-coding the predator to pursue prey not in schools, and to pursue a random prey at each time step if all prey were properly schooled, we were able to motivate the prey to form schools of three or more and to move together to avoid the predator. That the neural network was able to learn to school quickly, in some cases with only one hidden node, suggests that this task was much simpler than we initially thought.

Our findings support the work of Ward et. al. and other research that has found that neural networks can be used to train robots to exhibit schooling or flocking behavior[5]. Furthermore, our results are similar to those of Olson et al.[2]—we have found that a predator who becomes confused when prey travel in schools is sufficient motivation for said prey to adapt schooling behavior. We found that this schooling behavior does not require leader or follower prey, but only collective cooperation. This is demonstrated by the fact that the prey were able to move together, in a seemingly arbitrary direction, when being approached by the predator.

Directions for further work could include adding more constraints to the movement of the prey, such as prohibiting them from occupying the same space. We could also attempt to replicate the experiment in three dimensions in order to observe three-dimensional flocking patterns. After all, real fish move in three dimensions. Last, we could alter the experiment by having the predator be controlled by an artifical neural network instead of being hard-coded. The predator and prey would then co-evolve strategies for finding prey and avoiding the predator, respectively. This would be more similar to the work of Olson et. al.[2].

References

[1] Johnny O'Callaghan. Why do fish live in schools? http://www.howitworksdaily.com/environment/why-do-fish-live-in-schools, 2011. Accessed: 2014-05-01.

[2] Randal S. Olson, Arend Hintze, Fred C. Dyer, David B. Knoester, and Christoph Adami. Predator confusion is sufficient to evolve swarming behavior. *Journal of the Royal Society Interface*, 10:1–7, 2013.

[3] Francesco Pugliese and Davide Marocco. The relationship between flocking behaviour and the emergence of leadership. *Advances in Artificial Life, European Conference on Artificial Life*, pages 679–685, 2013.

[4] Kenneth Stanley. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21, 2004.

[5] Christopher R. Ward, Fernand Gobet, and Graham Kendall. Evolving collective behavior in an artificial ecology. *Artificial Life*, 7:191–209, 2001.