Evolution of Herding Behavior in Artificial Animals

Gregory M. Werner
Michael G. Dyer

AI Lab/Computer Science Department
University of California
Los Angeles, CA, 90024

Abstract

We have created a simulated world ("BioLand") designed to support experiments on the evolution of cooperation, competition, and communication.

In this particular experiment we have simulated the evolution of herding behavior in prey animals. We placed a population of simulated prey animals into an environment with a population of their predators. The behavior of each of the animals is controlled by a neural network architecture specified by its individual genome. We have allowed these populations to evolve through interaction over time and have observed the evolution of neural networks that produce herding behavior. The prey animals evolve to congregate in herds, for the protection it provides from predators, as well as the help it provides in finding food and mates. An interesting evolutionary pathway is seen to this herding, from aggregation, to staying nearby other animals for mating opportunities, to using herding for safety and food finding.

1 Introduction

We are pursuing a bottom-up exploration of the evolution of cooperative behavior and communication.

Our work has involved building a simple model of the world ("BioLand") that contains information relevant to the day-to-day life of simple animals (termed "biots"). This environment is then populated with biots whose behavior evolves to best survive in interaction in the environment. If the relevant features of the real world are included in the simulation, the biots should evolve cooperative skills similar to those seen in nature.

Each biot is controlled by a neural network, specified by its own genome. By examining those neural networks that evolve, we might be able to gain some insights into how animal brains might accomplish similar cooperation tasks. In addition, we believe that the task of constructing BioLand simulations will reveal aspects of what the actual constraints, goals, etc. of cooperation might be like in the real world.

2 Herding, Flocking, and Schooling

Many types of animals seek the company of others. There are several benefits to be gained by gathering into groups with others of the same species, including safety from predators, access to mates, and help in finding food. The advantages discussed here for animals that herd are true in large part, for example, with flocking birds and schooling fish.

Above all, herding provides safety for the individual. It increases the effective vigilance of the individual, can confuse or intimidate a predator, and can be used to provide cover where none exists. It also makes it less likely for predators to find prey, perhaps limiting the predators' numbers.

When animals are feeding, they must occasionally stop and look around to avoid being eaten by a predator. If an animal can tell if another in the herd senses danger, it no longer has to spend as much time looking around.

In addition, simply having more individuals looking makes it less likely that danger will be overlooked. More eyes and noses working together increase the chance that a predator will be located before it is within striking distance.

If a predator approaches the herd, the large numbers of animals in the herd can be used to fight off the predator. Circular defensive formations are used by some animals, and others may attack or mob the predator together. Musk oxen form a circle with their horns pointed out to deter predators, and baboons will mob leopards to chase them off [7].

The numbers of similar appearing possible targets also can confuse a predator. Once the predator tries to attack an individual, other animals moving around nearby can distract the predator, perhaps making it change its target animal repeatedly instead of pursuing one animal.

Finally, each animal may attempt to hide behind another, less fortunate herd-mate. This is known as selfish herding. Basically it is using another animal as a shield. Each animal in a herd may struggle to stay in the relatively safe center of the herd at the expense of the others. This hiding behind other animals clearly can keep a herd together. This desire to remain in the center of the herd must be balanced with the greater access to resources often found at the periphery of the herd.

Herding also provides increased feeding opportunities. Just as many animals watching for danger is better than one, a parallel search for food is better for each individual involved, given some kinds of food clustering.

If food is found in patches that can be shared, but where the patches themselves are somewhat difficult to find, there is pressure to cooperate in finding the patches.

Herding also brings together prospective mates, allows copying of behavior, and aids navigation.

3 Design and Goals

We hope to recreate an evolutionary orgin of herding. We are trying to simulate the evolution of herding behavior in animals that do not originally herd.

We hope to evolve herding in simulated prey animals, possibly with danger and food signals to help coordinate the herd. This could illustrate a possible evolutionary path to this cooperative behavior.

We also hope to evolve group hunting in simulated predators.

To these ends, we have set up an initial 2-D environment with 2 species of animals, which we call "deer", and "wolves." There are 2 kinds of nonanimal objects in the environment as well: "plants", and "trees."

Deer increase their energy level by eating plants. Wolves increase their energy level by eating deer. This energy is used up by the animals' baseline metabolism, and there is an energy cost for each action a biot makes. Mating, eating, and producing sounds all have specific costs, and movement uses more energy the faster a biot moves. Each of these values is a parameter that must be set reasonably to evolve interesting behavior (Tables 1 and 2).

New plants are added to the environment at a constant rate to provide energy for the animals. They are placed near other existing plants so that clumps

action	cost for deer
move	$(speed/maximumspeed)^2$
mate successfully	-200
mate unsuccessfully	-1
eat succesfully	+100
eat unsuccesfully	-1
make sound	-1
exist	-1
be born	+400

Table 1: Costs and gains for deer actions.

action	cost for wolves
move	$(speed/maximumspeed)^2$
mate successfully	-200
mate unsuccessfully	-1

eat succesfully	+energy of prey
eat unsuccesfully	-1
make sound	-1
exist	-1
be born	+400

Table 2: Costs and gains for wolf actions.

of plants grow. Trees are positioned at the beginning of the simulation, and cannot be eaten or moved. They are obstacles in this environment and can allow the possibility of more complex future biots forming simple maps of the environment.

To reproduce, an animal has to have a high enough energy level to be fertile, and must find another animal willing to mate.

Each type of object has a specific (non-volitional) sound or smell associated with it that diffuses away from its location, and can potentially be sensed by a biot. The strength of this smell decays exponentially with distance. The (non-volitional) sound produced by animals is proportional to their speed. Nonanimal entities (e.g. trees) produce smells of constant strength (Figure 1).

Mating and eating behaviors also have characteristic (non-volitional) sounds that can be heard by other animals. This allow other biots to make use of the information about what another biot is doing.

In addition, animals can voluntarily produce sounds of varying "frequency" and "loudness." Louder sounds propagate farther than do softer sounds. Different frequencies allow different sounds and sound combinations to be used and discriminated. These sounds can be used as signals to other animals.

This environment should contain enough features of the real world to allow herding to evolve.

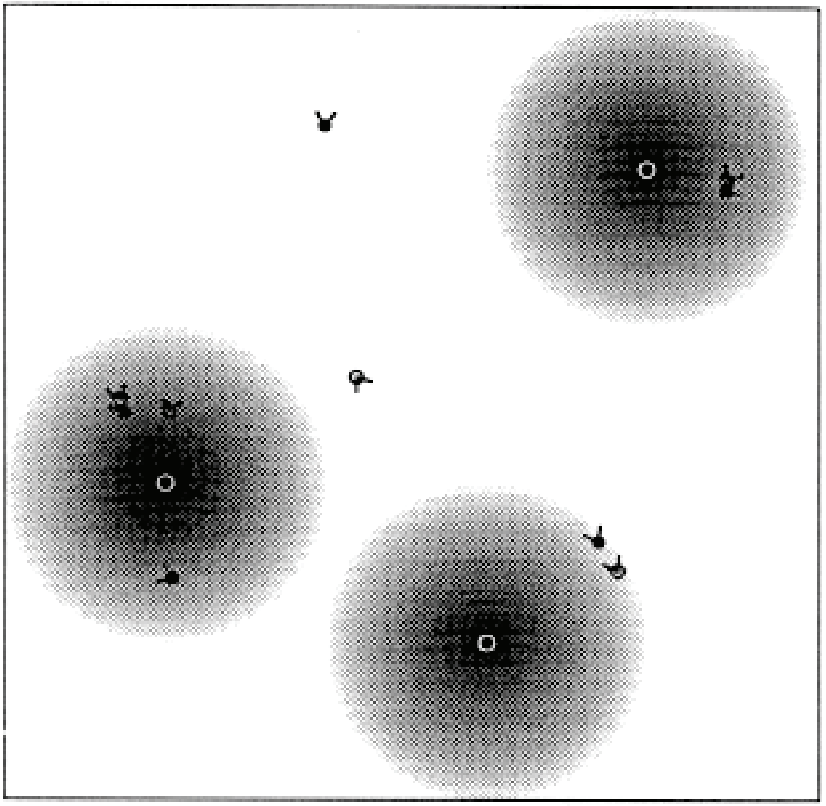

Figure 1:
Each object in BioLand produces a number of gradients.
Only plant smell gradients are shown here. Sensory neurons on
the ends of the animals' "antenae" fire proportionally to the
log of the gradient strength for which they are sensitive.

4 Modeling Biots

The behavior of each biot is determined by a neural network. Each biot perceives its environment in a species-specific way, then performs an action in response to its inputs. Biots have sensors on both sides of their bodies which sense the strength of each of the smell and sound gradients. There are specific sensory neurons for each smell and for each sound "frequency." The activation of these sensory neurons is proportional to the logarithm of the gradient strength. This is important because the exponential decay of sounds and smells creates gradient strengths over several orders of magnitude. In addition to sensory and motor neurons, each biot contains hidden neurons (3 in the current model) and higher order gating (axoaxonal) connections (Figure 3).

The neural networks that determine the behavior of the animals are themselves evolved. The sensory neurons and motor neurons of each species are predetermined before the experiment, and connections, hidden units, and axoaxonal connections are allowed to evolve. This allows arbitrary internal architectures to be specified given specific inputs and output capabilities.

Each animal has a genome that encodes the architecture of the neural network for that particular

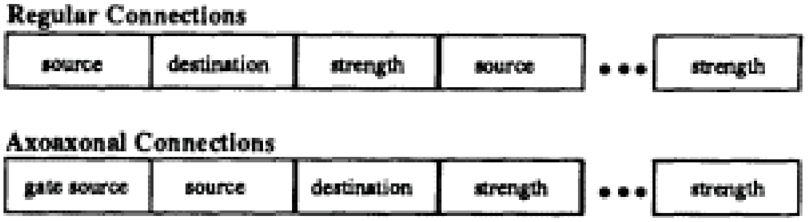

Figure 2:
The genome for a biot is split into 2 parts. The first part encodes the regular connections made between two neurons. Bits along the string encode first the source neuron (where the connection originates), then the destination, and finally the strength of that connection. The second part of the genome encodes axoaxonal connections. The first several bits here encode which neuron is the source of the gating connection. The next bits encode which connection is actually gated. Finally the gating strength is encoded.

animal. When animals mate, the genomes of the two parents are combined using crossover and mutation operations at the bit level [5] [6] and a new genome for the offspring is produced. This genome is then translated into the neural network that controls the new animal. Parameters are set so that each animal has one crossover on average, and has a 10 percent chance of having a mututation. Because of the coding used, each biot can effectively have a genome of any length between zero and 400 bits.

The scheme we use for encoding a neural network on a genome is a modification of that proposed by Collins [4] who represents a genome as containing a list of triples. Each triple encodes a single connection between 2 neurons. The triple contains the identification (ID) of the source neuron, the ID of the target neuron, and the strength of the connection between them. We have found that this scheme rarely evolves a hidden layer of neurons in between the input and output neurons. Since resulting networks are only 2 layers, and such networks are limited in their functionality, there are many behaviors that are not achievable.

To get around this problem, we have added a separate genome that encodes axoaxonal connections in the network. These are multiplicative connections that gate normal neuron-to-neuron connections, thus adding considerable computational power to a 2 layer network [3]. We encode each axoaxonal connection as a tuple with 4 elements. Each 4-tuple contains the ID of the source (gating) neuron, the IDs of the two neurons whose connection is being gated, and the strength of the gating synapse (Figure 2).

This gating can be used a number of ways to improve biots' behavior. A neuron detecting the presence of a predator can gate the connection between food and eating, for example. This would keep the biot from stopping to eat when it was in danger.

5 Experiment

We placed 8,000 prey animals and 8,000 predator animals into a 1000 by 1000 biot-length, toroidal environment.

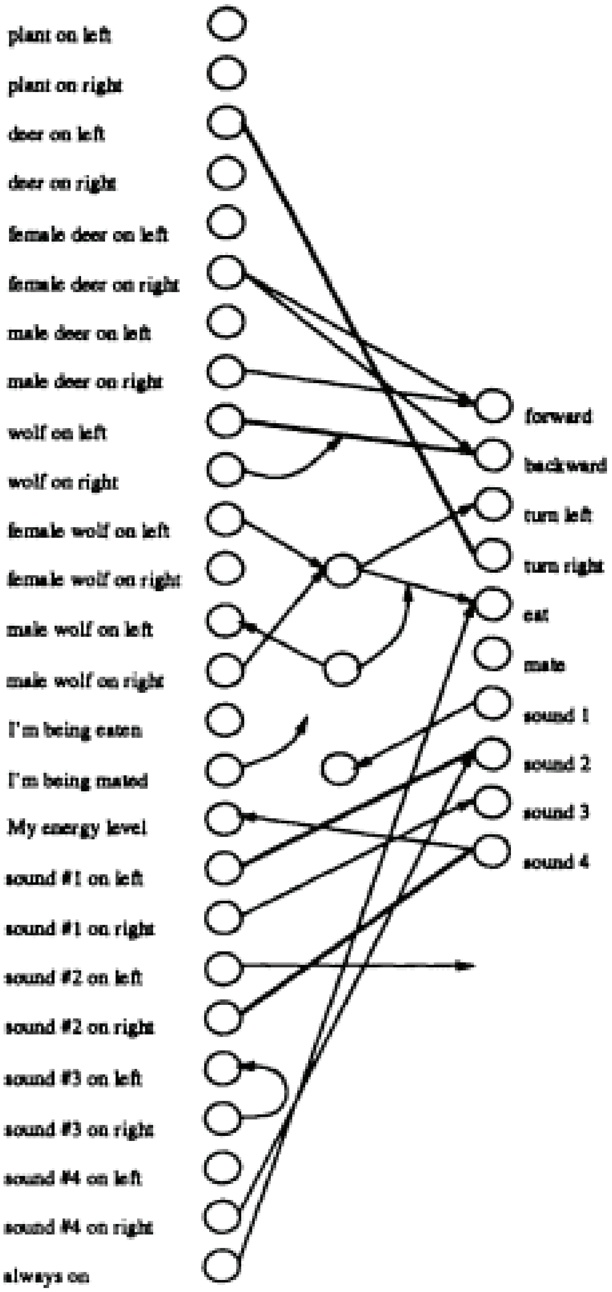

Figure 3:
A random neural network from the initial population. Thickness indicates strength of the connection. Axoaxonal connections (curved arrows) act to gate other connections. Note that some connections are "junk"in the sense that they are fragmentary or not interpretable (e.g. connections from input to input layer).

Each animal had a random beginning genome. Each genome can encode up to 25 regular, and 25 axoaxonal connections. The total number of bits in the genome was 400, so that there are 2^{400} total possible biot neural architectures.

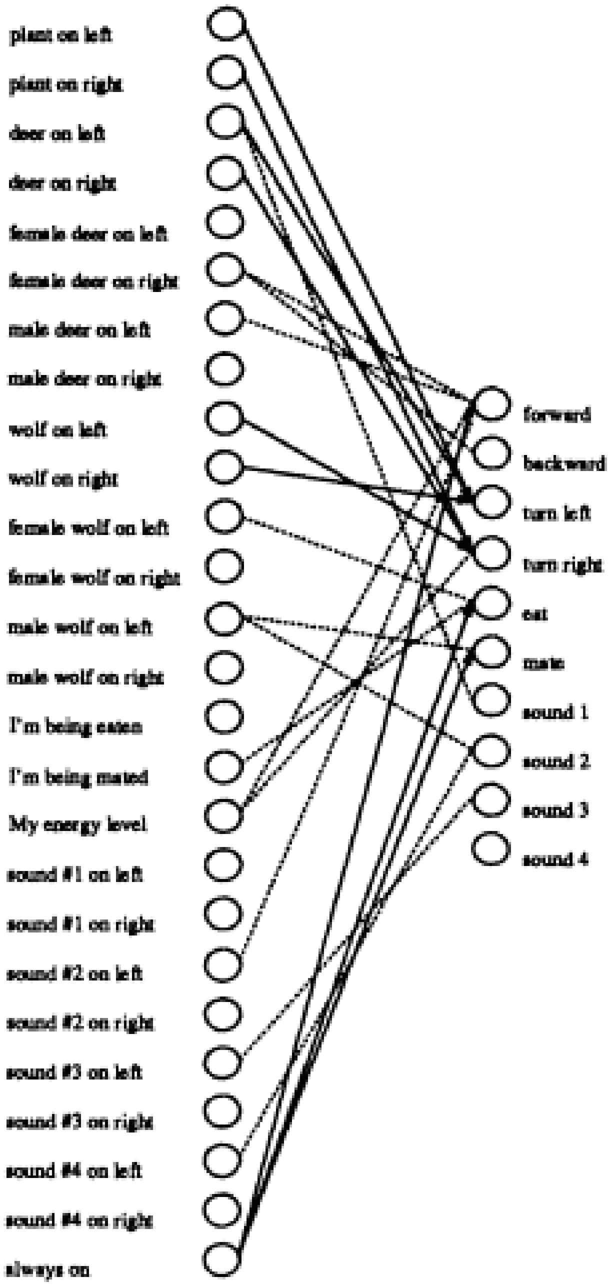

Figure 4:
The neural network of an evolved female deer after 80 generations.
Axoaxonal connections have been ommitted for clarity. Solid lines are
shown for connections which are common in female deer and clearly useful.
Dotted lines indicate connections that are less common and presumably less useful.

The behavior of the populations was checked occasionally throughout the run, and after approximately every 20 generations, a sample of the current genomes was taken. This sample was then analyzed to determine how the animals were implementing the behavior seen. Since there are not explicit generations in the simulation (because the animals mate and produce offspring at different times), the "generation length" is somewhat inexact. Every time the population produced a number of offspring equal to its average population size, we considered it another generation. Initial generations have a large number of sense/move cycles. Subsequent generation times shorten as biots become better at finding mates.

6 Results and Analysis

Early in the run, the behavior of the biots is random. Most of the biots die off because they cannot find food. The population shrinks down to about 400 biots. These have enough capability to eat and mate. Viable biots appear within the very first generation because all that is needed for eating/mating is that there exist some connection from the appropriate olfactory neurons to the appropriate motor neurons. For example, if a male biot moves toward the scent of a female biot and continuously executes the mating motor neuron (and the female does the same) then mating will occur. This behavior only requires the existence of a few connections. Later evolutionary pressures can then operate on these genotypes to improve mating efficiency (e.g. only firing the mating motor neuron when right next to a mate).

After about 30 generations, the predators have evolved to home in on the smell of prey, and the prey have evolved to move away from predators. This is accomplished by simply turning toward the side of the body in which the smell gradient is stronger in the case of the predators. Prey animals simply turn away from the stronger predator smell. In addition, the prey animals have evolved to home in on plants.

Once these behaviors are common in the population, we see that the animals tend to clump together. While they move in groups, this is not yet herding, but aggregation, because the animals are not seeking each other's company. They are simply pushed together by environmental features (i.e. food sources, places away from predators). Since the food is clumped together, the animals tend to clump together.

After 80-100 generations, we see that there evolves a tendency to move toward other animals of the same species. We can tell this by noting that neural connections that produce this behavior have arisen in the biots' neural networks (Figures 4 and 5).

At this point we see that most of the animals converge into small herds which are constantly splitting up and reforming as the groups interact with each other (Figures 6, 7, and 8).

Herding probably starts as a way of locating a mate. However, in this simulation there are at least 3 distinct pressures on the animals to do this. Once it occurs, the grouping of animals is true herding, because they are seeking to be in groups.

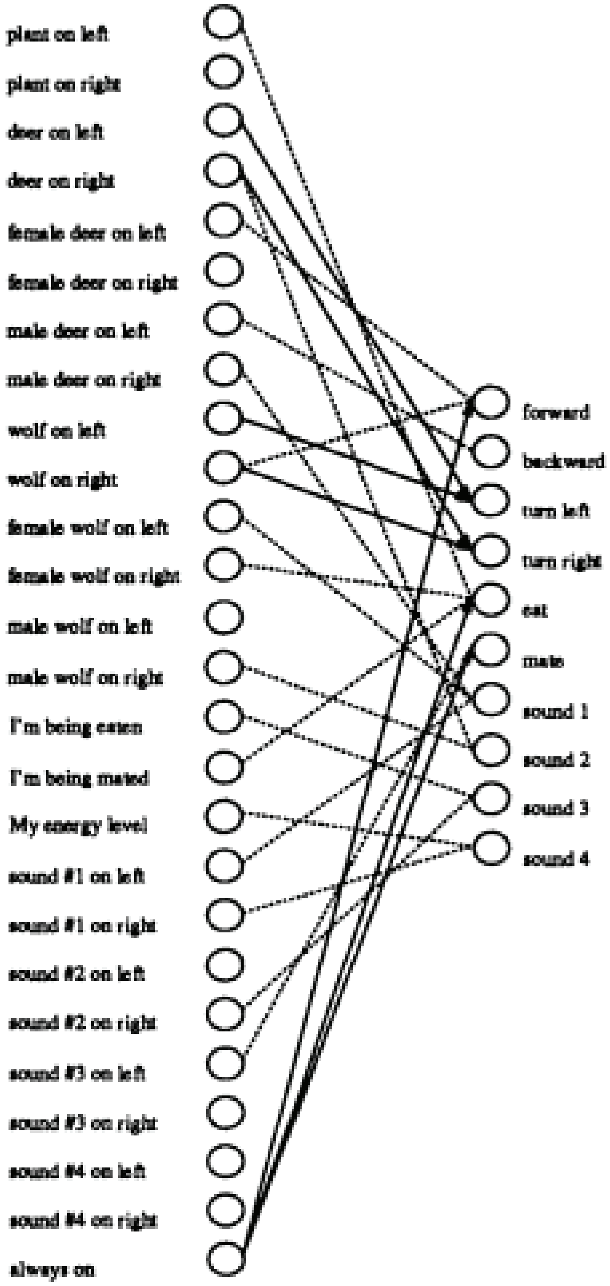

Figure 5:
The neural network of a male wolf after 80 generations.
Axoaxonal connections have been ommitted for clarity.

The first environmental pressure to herd is the pressure to find a mate. While aggregation provides mating opportunities, a biot that actually moves toward prospective mates clearly has an advantage over those who do not.

Second, there is pressure from predators. The prey "confusion effect" [7] is seen in these animals. A predator has a more difficult time in homing in on one prey animal than it does on more than one. We tested this case by placing a single predator animal into a test environment with one or more prey animals. We found that the predator animal took longer to capture a prey animal when the predator was in a group of prey than it did for a single animal. This was because the smell gradients produced by the prey interfered with each other so that there was not a single gradient to follow directly. The slope of the gradient was smaller, causing the predator to turn less quickly toward any one of the animals.

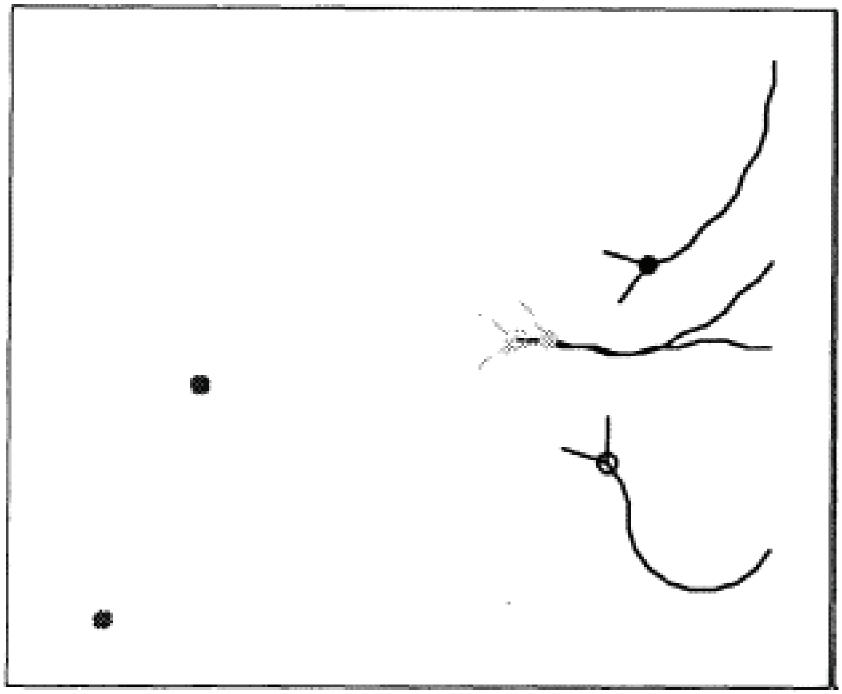

Figure 6:
Actual graphical output of lifelines of four biots in Bioland. Deer are shown in gray, wolves in black. Males are filled circles, females outlined. The "antennae" on the biots indicate where their sensors are located. Plants are black circles without antennae. Lines trailing the biots show their path. Note that the two deer have moved together, their paths intertwined and two wolves have begun to pursue them. Elapsed time here is 12 steps.

Third, there is pressure to find food. A group of animals can sense a larger area of the environment than can one.

If a animal senses food that the others cannot, it will move towards the food, slightly drawing the group with it. The rest of the group will then get to the point where they can sense the food as well, and the whole group will move to the food source. Because we made the food in this environment grow in patches, there is an advantage to cooperatively finding food.

Why is this model as successful as it is? We think there are several reasons:

1. There is plenty of food in the environment to begin with (8000 plants). Furthermore, the massive death of the initial biot population leaves plenty of food for those who have survived.

2. Biots do not die through aging. Therefore, any biot who has evolved to eat can live almost indefinitely, thus increasing the likelihood of mating.

3. The input neuron, termed "always on," represents a kind of internal "drive." Without this neuron, biots are completely reactive which is not realistic. A biot who only moves when it senses food, for example, will become immobile and therefore trapped if there is no gradient around it of sufficient intensity.

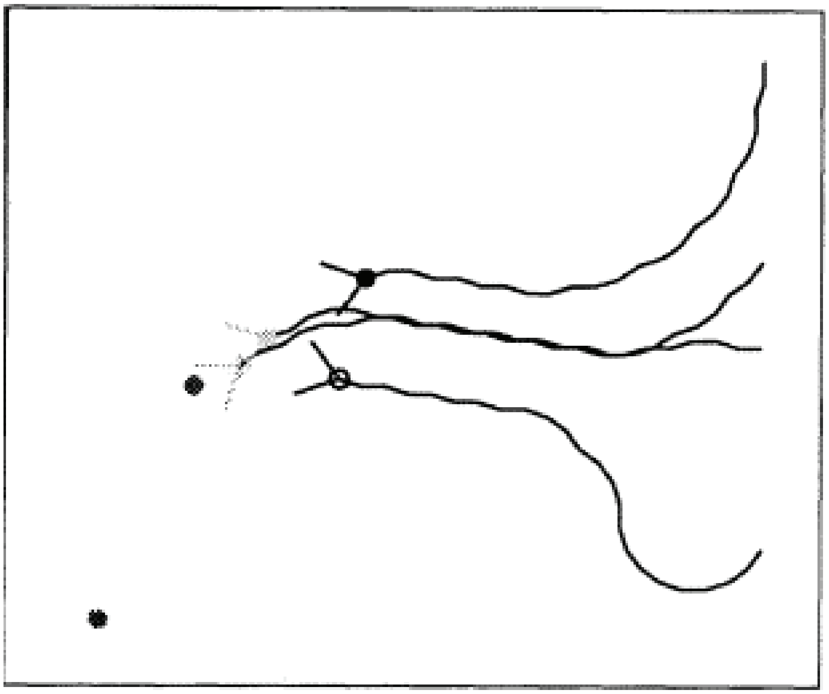

Figure 7:
The wolves have gained on the deer over the past
12 time steps. The deer are moving toward the plant
to eat it, while at the same time avoiding the wolves.

4. The neural net evolutionary space for this task is rather smooth. That is, biots don't have to have perfect connections to survive. For example, even if the "always on" neuron is not connected, there may be enough connections from other sensors to successfully move the biot toward food and mates.

7 Current Implementation and Future Work

BioLand is implemented on a 16k processor CM2 Connection Machine whose massive SIMD style parallelism is well suited for such tasks. It is written in C* and it took several days to run the experiment described here. Simulating the gradients in the environment is the major bottleneck.

Just as herding among predators has evolved, we would like to evolve group hunting by predators. We hypothesize that the following circumstances will need to be in effect: (a) prey must be able to outrun predators, (b) prey must be able to metabolically outlast predators, and (c) prey must provide more food than a single predator can consume. Currently, prey and predators move within the same velocity ranges and when a predator "kills" a prey, the predator's metabolism is increased by an amount equal to the prey's metabolism. There is nothing like a "carcass" being left as additional food in the environment. If prey are larger and faster than predators, then predators who hunt together should evolve. Group hunting, for example, allows one hunter to rest while the others continually force the prey to "burn" metabolism and thus tire out. The resulting kill can then supply food for the entire "pack".

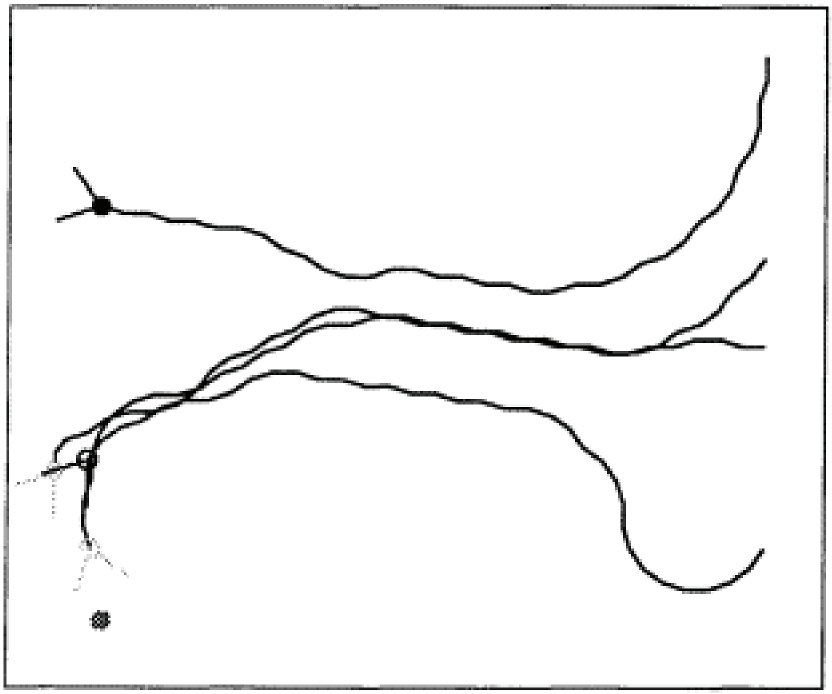

Figure 8:
In the next 12 time steps, several things have happened.
One of the deer has eaten the plant along their path. The male
wolf has moved off to pursue something off-screen. The female
wolf is about to consume the male deer. And the female deer
has broken away to consume a meal of a nearby plant.

An interesting additional direction is to explore under what circumstance volitional signaling might evolve among predators who hunt in groups. If predators have weak olfaction but good audition, then they might evolve to signal to each other the location of the prey as it runs out of sensory range.

8 Conclusions

We have shown how herding behavior could have arisen in previously non-herding animals. Pressure to locate a mate and the clustering of food sources could have originally brought animals together simply as a side-effect of the environment. Subsequent advantages for prey would then arise, namely: (a) confusing predators by remaining in groups and (b) increasing the likelihood of finding more distant food clumps by following members who have already sensed the food and are following its gradient. These advantages would generate selection pressures for explicit herding to have spread throughout the evolving population. Once herding is in place, more coordinated group behaviors could then evolve (e.g. group defense formations and explicit signaling between herd members).

9 Acknowledgments

This research is supported in part by NSF grant IRI-910730 to the second author.

The CM-2 Connection Machine was acquired through NSF equipment centers grant BBS-87-14206, maintained through NSF grant DIR-90-24251. The CM-2 is managed by the UCLA Cognitive Science Research Program.

Thanks also to Liane Gabora and Adam King for insightful comments.

References

[1] John A. Alcock. *Animal Behavior: An Evolutionary Approach.* Sinauer Associates, Inc., Sunderland, MA, 1984.

[2] John T. Bonner. *The Evolution of Culture in Animals.* Princeton University Press, Princeton, New Jersey, 1980.

[3] Valentino Braitenburg. *Vehicles.* The MIT Press, Cambridge MA, 1984.

[4] Robert J. Collins. *Studies in Artificial Evolution.* PhD thesis, UCLA, 1992.

[5] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Welsey, 1989.

[6] J. H. Holland, K. J. Holyoak, R. E. Nisbett, and P. A. Thagard. *Induction. Processes of Inference, Learning, and Discovery.* The MIT Press, Cambridge, MA, 1986.

[7] D. McFarland, editor. *The Oxford Companion to Animal Behavior.* Oxford University Press, Oxford, 1987.

[8] C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. In *Procedings of SIGGRAPH*, 1987.

[9] Gregory M. Werner and Michael G. Dyer. Evolution of communication in artificial organisms. In *Artificial Life II*, 1991.