

Artificial Intelligence Laboratory, Computer Science Department, UCLA, Los Angeles, CA 90024

Evolution of Communication in Artificial Organisms

34. Wittgenstein, Ludwig. *Philosophical Investigations*, 3rd edition, section 19. New York, NY: Macmillan, 1958.

A population of artificial organisms evolved simple communication protocols for mate finding. Female animals in our artificial environment had the ability to see males and to emit sounds. Male animals were blind, but could hear signals from females. Thus, the environment was designed to favor organisms that evolved to generate and interpret meaningful signals. Starting with random neural networks, the simulation resulted in a progression of generations that exhibit increasingly effective mate-finding strategies. In addition, a number of distinct subspecies, i.e., groups with different signaling protocols or "dialects," evolve and compete. These protocols become a behavioral barrier to mating that supports the formation of distinct subspecies. Experiments with physical barriers in the environment were also performed. A partially permeable barrier allows a separate subspecies to evolve and survive for indefinite periods of time, in spite of occasional migration and contact from members of other subspecies.

INTRODUCTION

It is our goal to explore the evolution of language from simple genetically controlled signaling to learned patterns of communication that support complex forms of social interaction by simulating environments in which these types of interaction can evolve. As a first step in this direction, we have been exploring the evolution of simple intraspecies signals that are genetically hard coded into the behavior of simulated animals. Such innate signals are commonly found in the animal kingdom.⁸

By simulating environments that exert some pressure to communicate, we believe that we can evolve animal-like communication systems in artificial organisms. As our environments become more complex, we hope to obtain progressively more interesting communication systems. As the animals themselves become more complex, gaining learning ability, for example, we hope to see the evolution of primitive language in our artificial life populations.

Primitive communication is common in many species of real animals. For example, many animals emit signals that communicate their internal states or emotions. Signals representing hunger, fear, anger, or readiness to mate are all common among animal species. Young birds cry for food, lions growl at one another at the site of a kill, frogs issue calls to attract mates, etc. Signals are also used, though less commonly, to communicate something about the state of the external environment to other animals. For example, the location of food or the type of danger present may be signalled. Baboons let the other members of their group know that they have found food so that it can be shared. They also have separate danger signals for the presence of hawks, snakes, and leopards.⁸ Such signals seem to have a symbolic component to them. For example, a particular sound literally comes to mean "snake."

This work can also be viewed from another perspective. Agents that work simultaneously and communicate to solve a common problem can be considered a distributed algorithm. We are seeking to find general cases in which distributed algorithms can be evolved to solve non-trivial tasks. Evolving distributed algorithms could help to solve several basic problems in Distributed Artificial Intelligence. (See Bond and Gasser³ for an overview.) Tasks for which simulated evolution can be useful may be: decomposition of problems into subproblems to be solved by independent agents, selection of communication protocols, and the actual methods of solving each subproblem.

EVOLVING COMMUNICATION

We believe that a number of general principles should be followed when setting up simulations to evolve communication among organisms in artificial environments.

First, there should not be direct pressure on the animals to communicate. Communication should arise as a solution to another problem that has to be solved by the population. It is trivial to set up an evaluation function that directly rewards

animals that communicate. This approach, however, does not provide the population with the flexibility to evolve a communication system that the experimenter does not expect. Creative solutions would not be rewarded by an evaluation function that was biased in this way. Animals should not be judged on how well they communicate, but on how well they solve the task at hand. In this way, one can determine how communication aids in tasks that would normally be faced by an evolving population.

Second, it is important to present the populations with natural tasks such as finding food, protecting young, and attracting mates. These are the kinds of tasks that animals and humans faced when creating their communication systems, and such tasks placed important constraints on the development of communication. Since communication plays such an important role in cooperation, tasks that encourage cooperation are good prospects for bringing about communication.

EVOLVING DIRECTIONAL MATING SIGNALS

The specific problem we selected for the population to solve is mate finding. This is a problem that does not require communication, but can be aided by it. A good solution to this problem combines the concurrent evolution of search and signaling strategies.

To put evolutionary pressure on the animals to communicate, we needed to design animals in an environment such that some animals would have information that other animals needed to know but were not capable of finding out for themselves. The animals with this valuable information would have to communicate it to the other animals. The relevant information in this particular simulation is the location of the female animals relative to the male animals. The males do not know the location of female animals, and must listen to directions from females in order to avoid a blind search for mates. To accomplish this, we made the males of our species blind. The females, which we made immobile to avoid having the females simply find the male animals, must produce signals that will guide the males to themselves.

IMPLEMENTATION

The environment for our animals is a simple toroidal grid, 200-by-200 squares. The sides of the grid are wrapped around to avoid having animals getting stuck on the boundaries of the grid. Each of the 40,000 locations in the environment can be empty or occupied by one animal. Typically, we place 800 male animals and 800 female animals into the environment. Therefore, 4% of the locations are occupied at all times.

Each animal in the population has a distinct genome which is interpreted to produce the neural network that controls its actions. Each of the genes of the

genome has an 8-bit integer value that corresponds to the connection strength or the bias of a unit in the neural network of the animal.

The neural architecture is a recurrent network in which all hidden units are completely interconnected and have feedback to themselves.⁷ Recurrent neural networks differ from feedforward neural nets¹² in that they allow the network to use information about their previous state when producing their current output. Since no learning will take place, the weights and thresholds have been changed to integer values for faster program execution. Weights and biases are simply integers between 127 and -127. All thresholds are zero. The individual locations on the genome encode the weights and biases for the network (Figure 1). Each animal's genome has the encoding for both a male animal and a female animal. The sex of the animal determines which part of the genome is interpreted to create the neural network of the animal.

The female animal is given an "eye" that can sense the location and orientation of animals nearby. Specifically, she can detect any male animal that is within two squares vertically, horizontally, and diagonally of her location. Thus, she sits at the center of a 5-by-5 "visual field." Each of the nearby locations and orientations is associated with an input node in the neural net that is turned on when a male is seen in the particular location and orientation to which the node is sensitive. The female animal has no sensors for seeing or hearing nearby females. Also, unlike males, female animals have no specific orientation.




FIGURE 1 Each gene codes a connection weight (or bias) ranging from -127 to 127. Each artificial neuron has a threshold of zero.

The location of a male, within range of the female, produces a pattern of activation on the female's input units ("eye"). The squares around the female are ordered "closest" to "farthest" from her. If more than one male is within her receptive field, only the male that is "closest" is detected. Activation flows through the recurrent net of the female and produces a pattern of activation on her output units. This pattern of activation is interpreted as a "sound" that is transmitted to all of the male animals within the visual field of the female.

The male animals are given an "ear" that can hear these signals produced by nearby females. If more than one female sends a signal to a male, he hears only the sound produced by the female who is closer to him. Yes are broken arbitrarily, but consistently.

The outputs of the male animal are interpreted as moves made by the animal. The four output units of a male correspond to moving forward, standing still, turning left, and turning right. The action corresponding to the output unit with the highest activation is taken.

At each time step, each female's input units are set according to the location of male animals within her visual field. The input values for each of the females are then propagated through the neural net to produce a new output. These outputs, the "sounds" produced by the females, are then fed into the inputs of nearby males. The input values for the males are then propagated to produce the new actions for the male animals. The males are then moved, rotated, or left still, and the time step is complete (Figure 2).

When a male finds a female (moves onto the same grid location that she is on), the animals mate and produce two offspring, a male and a female. The parents' genome, which encode their neural network brains, are combined using the standard genetic operations of crossover and mutation⁹ to produce the genome of the offspring. (The mutation rate used was 0.01% per gene. The crossover rate used was 2% per gene.) These offspring replace two old animals in the population and the parents are moved to new locations in the environment so that they can attempt to reproduce again. Simply leaving the animals in the same location would have allowed them to mate repeatedly. The animals removed are selected randomly from the population.

This process, which we call "XGA," is an extension of the typical genetic algorithm (GA) because gene strands are reproduced as soon as they prove their fitness, instead of being compared and reproduced at fixed time intervals. The genetic algorithm⁹ traditionally has discrete generations where each of the members of the population is simultaneously judged and possibly paired off with a mate. However, it is important in evolution of language to allow inter-generational communication. For this reason, and for greater realism in simulations, we have made all reproduction asynchronous, thus creating overlapping generations. In addition,

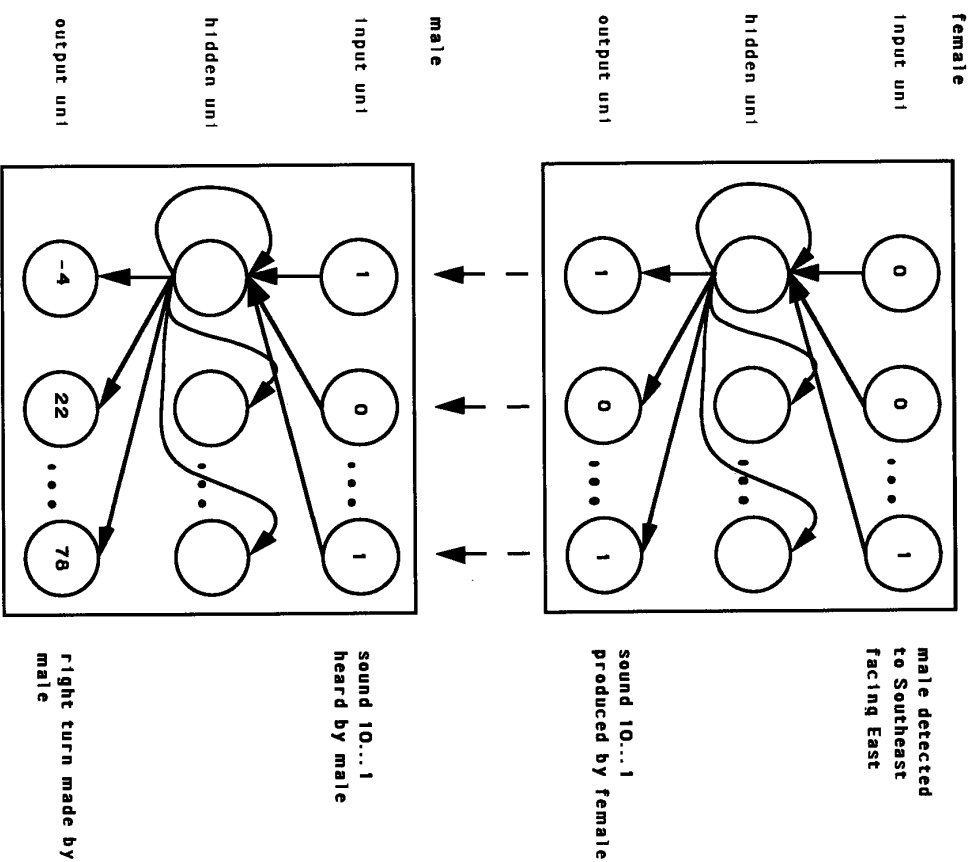


FIGURE 2 Female input units represent positions and orientations of males within females' "visual field." Males interpret female's output to aid in locating female.

in XGA, males must select one another directly, rather than being mated as a result of some fitness function unrelated to mate selection.

RATIONALE FOR DESIGN

To avoid having males simply learning to home in on a sound, the male was not given the ability to localize the sound source. He only hears the type of sound produced, but cannot determine the direction from which it came. This creates a situation in which the females must produce signals that have a meaning to the males, and the males must correctly interpret this meaning in order for the communication between the animals to be successful.

This formulation of the animals and environment was not meant to be an accurate simulation of any particular species of animal. It was adopted as a simple way of creating selective pressure on animals to communicate.

Within this environment, the best strategy for two animals to find each other is for the female to direct the male to herself. No search strategy employed by males can beat an efficient strategy of the female animals giving directions to males. Any search strategy employed by the males that doesn't include listening to directions from the female will not be as effective as the direction-following strategy. As the males are evolving a search strategy, females will have to evolve a corresponding signalling protocol for the males to follow.

EXPERIMENT

A simulation, with 13-bit wide inputs to the males, was run to determine whether a signalling protocol would evolve that could aid males in finding mates. In addition, the simulation was also run with the sounds produced by the females not copied into the inputs of the male animals. This provided an experimental control to determine whether the communication between males and females was actually used and useful to the males.

With three-bit wide outputs, the females are capable of producing $2^3 = 8$ distinct sounds; similarly, with three-bit wide inputs, the males are capable of recognizing eight different sounds. Since the output of the males consists of movements (left, right, still, and forward), we can interpret the sounds of the females as messages telling the males how to move.

However, the relationship between a specific sound uttered by a female and a move made by a male is arbitrary and depends completely on the genome of each (and its corresponding neural network). For example, one female may produce "011" when a male is one-right-turn-plus-one-forward-move away from her, while another female may produce the distinct sound "101" under identical circumstances. Likewise, upon hearing "011," one male may stand still; another, however, may turn left; yet another may go forward. What each male or female does, given its inputs, is completely determined by the weights on its neural network, which is completely specified by its genome (i.e., no learning occurs in these experiments).

The task is to co-evolve a population of males and females who agree on the same—albeit arbitrary—interpretation of the eight signals. No one "dialect" is *a priori* correct. More than one "sound" can be mapped onto the same motion. For example, a given male may turn right when hearing any seven of the eight possible

sounds and then move forward only when hearing the eighth. Although such a male can find a female, it will have a greater chance of getting stuck "spinning in place," than one who interprets more sounds as meaning "move forward." In any case, only one of the large number ($8^8 =$ over 16 million) of all possible making protocol dialects need co-evolve in order for females to successfully communicate with males.

While we may speak of females as "intending" to communicate, say, the message "turn right" to a male, there really are no intentions as such, since each animal's behavior is completely deterministic. Also, the interpretation of what a female message "means" can itself be problematic; for example, it may be the case that females are communicating to males messages more of the sort "hot" (i.e., "you are closer") and "cold" (i.e., "you are farther away"). In the following experiments, we describe the "meanings" of female messages in terms of the motions (right, left, forward, still) taken by the males.

RESULTS

During runs of the simulation, the behavior of the animals changed as improved mate-finding strategies were adopted. These changes in behavior occurred in the following stages.

1. Male animals wandered randomly and female animals signaled randomly. Since the animals started out with a random genome, they had neural networks with random connection strengths and biases. Therefore the population was full of male animals that moved erratically and females that emitted signals while oblivious to their surroundings (Table 1).

TABLE 1 Responses of randomly generated male neural nets to female signals are random (percentage of males making each response at Time = 100).

Signal	Move Forward	Turn Right	Turn Left	Stand Still
000	25	38	9	28
001	19	25	31	25
010	28	22	26	24
011	29	25	29	17
100	25	27	26	22
101	25	25	26	24
110	19	20	31	30
111	22	27	21	30

TABLE 2 Few males interpret a female signal to mean "stand still" (percentage of males making each response at Time = 5000).

Signal	Move Forward	Turn Right	Turn Left	Stand Still
000	74	20	5	1
001	81	12	7	0
010	67	18	14	1
011	79	11	19	1
100	80	7	12	0
101	75	11	14	0
110	56	21	23	0
111	70	14	16	0

2. Males that stood still became extinct. It is never a good strategy in this environment for a male to stand still. He should always move in order to cover the most ground and have the best chance of finding a female. Males that stand still lower their chance of finding a mate and therefore are selected against. Therefore very few animals interpreted any signal as meaning "stay where you are" (Table 2).
3. Males that usually go straight took over the population. Males that spent a large percentage of their time spinning in place were gradually replaced in the population by other males that spent more time covering new ground. Even though males may have evolved so that they would make appropriate moves when near some females, this ability was usually fatal when near a female that used a different signaling protocol. This is because the female's directions would steer him away from her, or very likely, would direct him to spin in place. For this reason, males evolved to simply ignore their inputs (Table 3). Moving in a straight line is a good search strategy because by avoiding turns, which cover no new ground, it covers the maximum amount of territory possible. Although this strategy covers only one row or column of the environment, the population is dense enough that on average there will be four females in each row or column of the environment. A male that follows this strategy will find a number of mates directly proportional to the density of the female population.
4. Males appeared that turn when in the same row or column as a female. At this point, females had evolved that produce a signal telling the males how to find them (Table 4). Females that did not use these signals, or that gave inappropriate ones, gradually became less common in the population. A male can maximize his chances of reproducing by going straight when not near a female, and by listening to a female's directions when close to one.

TABLE 3 Males ignore their inputs and simply move forward (percentage of males making each response at Time = 7500).

Signal	Move Forward	Turn Right	Turn Left	Stand Still
001	98	2	0	0
001	99	0	0	1
010	98	1	1	0
011	100	0	0	0
100	100	0	0	0
101	99	1	0	0
110	98	1	1	0
111	99	0	1	0

TABLE 4 Males evolve that interpret 101 as "turn left," 110 as "turn right," and the remaining patterns as "move forward" (percent of males making each response at Time = 15,000).

Signal	Move Forward	Turn Right	Turn Left	Stand Still
000	97	2	0	1
001	100	0	0	0
010	98	1	1	0
011	98	0	1	1
100	100	0	0	0
101	22	77	1	0
110	5	2	93	0
111	97	0	3	0

5. Female animals evolve to use the existing signals in more situations. Typically, the first use of a "turn" signal is in the case where a male is adjacent to a female and only needs to turn once to find her. Once males evolve to turn according to this signal, the females use it in more and more situations.

After about 50,000 time steps, females evolve so that they will signal males who happen to be on the same row or column as the female to turn towards the female. A cross-shaped area (Figure 3) appears in which males will be told to spin until they face the female. All other males sensed by the female will be signalled to go straight. One can see that this strategy is the quickest way for females to guide males to themselves.




FIGURE 3 Females evolve that tell males in cross-shaped (shaded) area to continue turning until the male is facing the female. Males entering the "receptive field" but outside of the cross-shaped area are told to go straight. Such males will eventually enter the shaded area and from there be guided in to the female.

In about half of the runs, only one signal for turning evolves. Since turning one way can be accomplished by turning the other direction three times, it is possible to do without a signal for one of the turns. Since it is less efficient, however, one would expect that eventually a signal representing a turn in the other direction would evolve.

COMPARISON WITH NON-COMMUNICATING ANIMALS

To test how much the ability to communicate helped these animals, it was compared to the control group in which the males could not hear the signals produced by the females (Figure 4).

Before time = 7500 the population in which the males ignored the females did better than the "listening" males. This is because some of the listening males followed bad directions from a female. These bad directions were the result of incompletely evolved signals. For example, many males got caught because they were repeatedly told by a female to turn in place. They eventually died and were replaced by offspring of a more successful male.




FIGURE 4 Animals with communication evolve to find mates in less than half the time required by animals lacking communication.

Between time = 7500 and time = 15000, the population with “listening” males reproduced more rapidly than the control population, but still below the rate possible using the best strategy that does not involve communication (i.e., moving in a straight line).

By time = 15000, the population with “listening” males reproduced more rapidly than could any population not employing communication.

Note that the control group never reached the maximum possible reproduction rate. Mutations away from the optimal strategy were common enough that about 25% of the animals would behave in non-optimal ways. Since many offspring were produced that were worse than their parents, the average time to find a mate was less than optimal for the population with deaf males. By simply travelling in a straight line, and thereby covering the maximum number of squares possible, the deaf males could bring their average time to find a mate as low as 50 moves. In practice, they never did better than finding a mate every 100 moves. “Listening” males, however, found females in an average time of 40 moves.

EXAMPLES OF ANIMALS USING EVOLVED SIGNALS

EXAMPLE 1

This is an example of the use of the evolved communication protocol from the particular simulation described above.

1. A male is just outside of visual and acoustic range of a female. He does not hear the signal being output by the female, nor does the female see him (Figure 5a).
2. The male has moved into the view of the female animal by moving straight. This was the optimal move given that he couldn't hear anything. At this point the female signals the male to “move forward” (Figure 5b).
3. The female signals the male to continue to “move forward” (Figure 5c).
4. The male has reached a square adjacent to the female, but will continue past her if he does not turn. Appropriately, the female changes her signal to one meaning “turn right” (Figure 5d).
5. Following the signal from the female, the male turns to his right. The female now changes her signal to one meaning “move forward” (Figure 5e).
6. The male then moves onto the square occupied by the female. Notice that the interpretation of a specific signal that evolves will be different in each simulation run. For example, 101 means “move forward” in Figure 5 while it means “turn right” in the run that produced Table 4. Although the specific signals may be different, the overall evolved protocol is the same.

EXAMPLE 2

This is an example of the use of a protocol evolved in a different run of the simulation. In this run, no signal for “turn right” has evolved.

1. A male is just outside of visual and acoustic range of a female (Figure 6a).
2. The male has moved into the view of the female animal by moving straight. The female signals the male to “move forward” (Figure 6b).
3. The female signals the male to continue to “move forward,” but uses a different signal with this same meaning (Figure 6c).
4. The male has reached a square in the same column as the female. The female lacks a signal that means “turn right,” so she begins signaling for a sequence of left turns that will cause the male to face her (Figure 6d).
5. The female again signals the male to turn left (Figure 6e).




FIGURE 5 a) Female's default signal is "move forward." b) Female signals "move forward." c) Female signals "turn right." d) Female signals "turn right." e) Female signals "move forward."




FIGURE 6 (see caption next page)

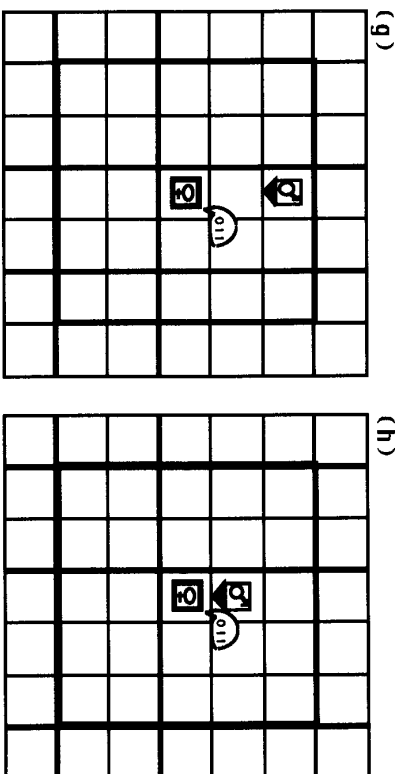


FIGURE 6 (continued) a) Male approaches "receptive field." b) Female signals "move forward." c) Female indicates "move forward" with another signal. d) Female signals "turn left" since she lacks a signal for "turn right." e) Female signals "turn left" again. f) Female signals "turn left" again. g) Female signals "move forward." h) Female signals "move forward" again.

6. The female signals the male to make a final left turn (Figure 6f).
7. Now that the male is facing her, the female changes her signal to one meaning "move forward" (Figure 6g).
8. The male is now adjacent to the female. Only one more "move forward" signal is required (Figure 6h).
9. The male then moves onto the square occupied by the female.

EVOLUTION OF DIALECTS

One would expect that in an environment such as this one, there may be many possible mappings of signals to meanings that could solve the task. Each of these mappings can be considered a different "language" or "dialect." To aid in viewing the evolution of these dialects, we both modified and simplified the experiment somewhat.

The animals were reprogrammed as simple pattern transducers (i.e., hidden layer and recurrent connections were removed). Each gene in the female animal encoded a signal that the female would emit when a male was at a specific location relative to her. Each gene in the male encoded his response to each possible signal

from the female. We also reduced to four (i.e., two bits) the number of possible signals that the female could produce. This way, we reduced the total number of types of male animals to $4^4 = 256$ (four possible inputs for each of four possible moves). This representation has the advantages of being much easier to analyze, evolving faster under genetic operations, and running faster.

RESULTS

A series of tables follow that show the evolution of signal responses by the males. Each position in a given table contains two numbers separated by a colon. The first number is the frequency (i.e., number of males) of that particular set of signal responses. The second four-digit number represents the signal responses themselves. The first of these digits represents the response of the male to signal #1 from the female. The second digit represents the response to signal #2, and so on. In these responses, 0 means stay still, 1 means go forward, 2 means turn left, and 3 means turn right. Therefore, the entry 7:0113 means that there are 7 animals that will stand still if they hear signal #1, move forward if they hear #2 or #3, and turn right if they hear #3. Each table contains 256 logically possible dialects, from 0000 (males that always stand still) in the upper left-hand corner to 3333 (males that always turn right) in the lower right-hand corner.

The particular run shown here is interesting because a non-optimal strategy eventually takes over the population. A protocol evolves that only allows "move forward" and "turn right." Apparently, *this dialect dominates because the males that employ it are "bilingual."* They respond correctly to signals from two other common dialects.

Frequency Matrix 1 (Table 5): At time = 0, most of the possible communication protocols are used by at least one male.

Frequency Matrix 2 (Table 6): At time = 8,000, most of the animals that stand still have died out. Several good protocols have become popular including 1311, 1211, 1321, 1112, and 1113. These good strategies allow males to move forward and to turn to find females when necessary. The protocol that eventually wins out, 1313, is used by only one animal.

Frequency Matrix 3 (Table 7): At time = 10,000, protocol 1313 starts to gain in population. We found that females that attracted males using protocol 1311 used only signals #1 and #2. Females that attracted males using 1113 used mainly signals #3 and #4. This allowed males that used protocol 1313 to follow the correct directions whenever it found either of these groups of females. In a sense, *the 1313 males are bilingual.*

TABLE 5 Frequency of Male Signal Responses (time = 0)

4:000	6:0001	5:0002	7:0003	4:0010	6:0011	4:0012	4:0013
12:0020	1:0021	3:0022	4:0023	8:0030	8:0031	2:0032	7:0033
5:0100	7:0101	7:0102	5:0103	7:0110	5:0111	3:0112	8:0113
2:0120	7:0121	6:0122	7:0123	3:0130	9:0131	6:0132	7:0133
8:0200	6:0201	4:0202	3:0203	11:0210	3:0211	4:0212	8:0231
1:0220	7:0221	6:0222	5:0223	8:0230	5:0231	8:0232	2:0233
8:0300	7:0301	2:0302	9:0303	4:0310	9:0311	9:0312	4:0313
10:0320	7:0321	3:0322	5:0323	12:0330	7:0331	6:0332	4:0333
8:1000	7:1021	6:1002	7:1003	7:1010	4:1011	7:1012	2:1013
7:1020	7:1021	0	7:1023	4:1030	8:1031	7:1032	8:1033
11:1100	4:1101	4:1102	6:1103	10:1110	3:1111	3:1112	8:1113
3:1120	5:1121	6:1122	4:1123	3:1130	8:1131	2:1132	7:1133
11:1200	4:1201	6:1202	5:1203	5:1210	6:1211	8:1212	5:1213
7:1220	5:1221	8:1222	8:1223	5:1230	4:1231	7:1232	5:1233
8:1300	4:1301	9:1302	7:1303	8:1310	11:1311	8:1312	6:1313
4:1320	3:1321	10:1322	6:1323	5:1330	2:1331	6:1332	7:1333
7:2000	9:2001	7:2002	11:2003	3:2010	9:2011	5:2012	6:2013
12:2020	8:2021	5:2022	7:2023	7:2030	6:2031	4:2112	10:2033
7:2100	5:2101	8:2102	3:2103	4:2110	7:2111	4:2112	5:2113
11:2120	11:2121	6:2122	6:2123	5:2130	8:2131	4:2132	9:2133
7:2200	4:2201	3:2202	10:2203	7:2210	8:2211	4:2212	7:2213
8:2220	7:2221	9:2222	11:2223	4:2230	5:2231	3:2232	5:2233
6:2300	8:2301	6:2302	7:2303	7:2310	7:2311	5:2312	6:2313
6:2320	4:2321	4:2322	4:2323	3:2330	3:2331	8:2332	3:2333
8:3000	6:3001	8:3002	4:3003	8:3010	4:3011	9:3012	5:3033
4:3020	8:3021	4:3022	9:3023	6:303	10:3031	6:3032	8:3113
9:3100	5:3101	7:3102	3:3103	14:3110	8:3111	5:3112	6:3133
7:3120	5:3121	5:3122	12:3123	3:3130	6:3131	6:3132	6:3133
4:3200	10:3201	6:3202	4:3203	12:3210	6:3211	10:3212	0
14:3220	9:3221	3:3222	12:3223	4:3230	6:3231	7:3232	7:3233
8:3300	4:3301	4:3302	9:3303	8:3310	8:3311	4:3312	3:3313
4:3320	6:3321	11:3322	6:3323	4:3330	5:3331	6:3332	6:3333

Frequency Matrix 4 (Table 8): At time = 12,000, protocol 1213 starts to appear more frequently. This is another "bilingual protocol" combining 1211 and 1113. Frequency Matrix 5 (Table 9): (18 rows of zeros have been removed.) At time = 14,000, protocol 1313 is now used by a sizeable part of the population. Some good protocols including 1321, 1111, and 1311 are starting to decline.

TABLE 6 Frequency of Male Signal Responses (time = 8000)

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	1:0101	0	0	0	0	0	0	1:0113	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	1:0222	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	1:0312	0
0	0	0	0	0	0	0	0	1:0313	0
0	0	0	0	0	0	0	0	1:0333	0
2:1000	0	0	0	0	0	0	0	1:1013	0
1:1020	0	0	0	0	0	0	0	73:1011	0
0	1:1101	0	0	1:1103	1:1110	0	2:1032	64:1033	0
0	131:1121	14:1122	2:1123	0	0	323:1211	0	24:1133	0
1:1200	0	0	0	0	0	0	0	5:1212	0
1:1220	4:1221	2:1222	12:1223	0	0	0	0	2:1233	0
4:1300	2:1301	1:1302	0	5:1310	547:1311	0	74:1312	1:1313	0
0	89:1321	0	1:1323	0	13:1331	0	0	1:1333	0
0	0	0	0	0	0	0	0	0	0
0	1:2021	0	1:2023	0	0	0	0	0	0
0	0	0	0	0	0	0	0	1:2113	0
0	1:2121	0	1:2123	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	1:3002	0	0	0	0	0	1:2332	0
0	0	0	1:3023	0	0	0	0	0	0
1:3100	0	0	0	0	0	0	0	0	0
1:3120	1:3121	0	0	1:3130	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
1:3220	0	0	0	0	0	0	0	0	0
0	0	1:3302	0	0	0	0	0	0	0
0	0	1:3322	0	0	0	1:3331	0	0	0

Frequency Matrix 6 (Table 10): At time = 16,000, protocol 1213 has surged ahead of 1313. The two protocols that it combines are now more common than those that 1313 combines. Frequency Matrix 7 (Table 11): At time = 20,000, protocol 1313 has become much more common than 1213, even though it is a less successful strategy.

between organisms. We would like to create environments that pressure a population into signalling internal states and intentions among the members, as well as letting one another know about states and events in the environment. We want to keep the tasks for our populations “natural”—similar to the selective pressures that actually brought about the ability to communicate to living systems on earth. We are also examining tasks that require more interesting interactions between animals to accomplish a common goal and that require the individuals to use state information when deciding on what actions and/or signals to produce.

IMPROVING THE MODEL

In order to achieve these research goals we are currently making a number of improvements to our model. First, the physics of the environment is currently very simple. The “sounds” produced by animals are simply copied to other animals nearby. These sounds have no direction or intensity, which would carry a large amount of information. Nor can an animal hear more than one sound at a time. The “vision” possessed by the female animals is also impoverished. The female is only able to detect the presence of a male animal within a very small area. “Vision” in the current simulation also does not take into account factors such as closer objects obscuring those farther away. In later experiments, we will eliminate these flaws. We plan to add sensors onto our animals that can detect the direction of the sound source and the intensity of the sound. Also, the eyes of our animals will no longer be able to see through objects.

The environment is relatively barren in the work described in this paper. The animals have the entire environment to themselves. This will change, as we plan to add a number of new kinds of objects to the environment, such as plants, rocks, pools of water, recognizable offspring, and other species of animals. With these new objects we hope to create a large number of new and realistic tasks for populations to solve.

Metabolism is a key feature which is also lacking in this particular implementation. Including it will force our animals to find food to stay alive (an interesting task by itself), and deaths by starvation will provide a natural way to keep the population at a reasonable size.

Avoiding predators is a very common task for real animals, so we plan to add predatory species to our simulations, hopefully to evolve both predator avoidance and possibly group hunting behavior.

The animals in our current system have a very small number of possible actions. We are adding a few more primitive actions to the animals’ repertoire, including the ability to grasp and release objects in the environment, and making as an action to be selected (versus simply occurring automatically as the result of cohabitating the same square).

The random placement of offspring is another flaw in the current model. By placing offspring near their parents in future models, we will avoid mixing up the

animals in the environment. We believe that this, along with the ability to recognize one’s offspring, may encourage both speciation and altruism toward kin.

The use of a direct mapping from genome to neural network connection strengths and biases has several major flaws. First, it is clearly biologically implausible. Secondly, it is impossible to use a direct mapping for large neural networks because of the exponentially increasing number of connections contained in them. Finally, this representation does not seem to work well with the genetic algorithm. Experiments with more complex environments and animals have shown that it is extremely difficult to evolve larger neural nets using this simple mapping. We believe that the genome-to-neural net mapping and representation issues pose major research challenges.^{5,6}

Lastly, there is a learned component to communication in addition to innate signals. Learning allows the more rapid creation and acceptance of new signals among a population. A new signal can be created through the invention of a new sound. We plan to add a form of unsupervised learning to our organisms which we hope will allow forms of communication to evolve with both innate and learned features.

BIOLOGICAL ISSUES

It is interesting to note that the use of acoustic signals in nature can constitute a behavioral barrier to mating—one that can serve as a basis for later speciation. For example, in nature there are several species of frogs that can produce viable offspring with members of other species. However, the frogs are considered distinct species because the calls made by the frogs insure that such mating almost never happens. The calls made by male frogs to attract females are only pursued by females that are of the same species.^{1,2}

We believe that communication protocols could provide a natural way of establishing genetic barriers that spontaneously emerge. This could be useful in exploring a wide variety of biological problems in the origin and maintenance of distinct species in an ecological setting, including kin selection, altruism towards kin, genetic drift, gene flow, mimicry, and parasitism.

CONCLUSIONS

We have shown that it is possible to evolve organisms that communicate to solve a simple make-finding task. The ability to produce appropriate signals for stimuli co-evolved with the ability to take appropriate actions upon receiving each signal. No explicit pressure on the artificial animals was required to cause this communication

to develop. It simply arose as part of a cooperative solution to the make-finding problem—as it appears to have in nature.

We have also shown that it is possible to vary the genetic algorithm so that the phenotype (neural network) produced by each genome is capable of finding a mate. This variant algorithm, dubbed "XGA," is more realistic than the standard genetic algorithm in which each genome is scored by some evaluation function at fixed intervals and randomly paired off with a similarly scoring animal. The XGA algorithm can be used without losing the ability to evolve useful phenotypes.

Finally, we have shown that subspecies or distinct "dialects" do evolve and compete. Dialects that are "bilingual" (i.e., correctly interpret several signalling protocols) have an increased chance of dominating in the long run. Physical barriers, however, allow distinct dialects to survive indefinitely in different regions, even in cases where partial permeability of a barrier permits some migration across regions and therefore contact between distinct subspecies.

ACKNOWLEDGMENTS

This work was supported in part by a W. M. Keck Foundation grant. Thanks to Rob Collins, David Jefferson, Adam King, Richard Korf, and Charles Taylor for their insightful comments on earlier drafts of this paper. Thanks also to Chris Langton and several anonymous reviewers for their suggestions on an earlier draft of this paper. A special thanks to Ruth Liow for her help in preparing this document.

REFERENCES

- Blair, W. Frank. "Isolating Mechanisms and Interspecies Interactions in Anuran Amphibians." *Qtr. Rev. of Biol.* **39** (1964): 334-344.
- Blair, W. Frank. "Mating Call in the Speciation of Anuran Amphibians." *Am. Naturalist* **92** (1958): 27-51.
- Bond, Alan H., and Les Gasser, eds. *Readings in Distributed Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann, 1988.
- Bonner, John T. *The Evolution of Culture in Animals*. Princeton, NJ: Princeton University Press, 1980.
- Collins, Robert J., and David R. Jefferson. "AntFarm: Towards Simulated Evolution." This volume.
- Collins, Robert J., and David R. Jefferson. "Representations for Artificial Organisms." In *Proceedings of Simulation of Adaptive Behavior*, eds. Jean-Arcady Meyer and Steward Wilson. Cambridge, MA: MIT Press/Bradford Books, in press.
- Elman, Jeffrey L. "Finding Structure in Time." Technical Report CRL 8801, Center for Research in Language, University of California, San Diego, 1988.
- Frings, Hubert, and Mable Frings. *Animal Communication*. Norman, OK: University of Oklahoma Press, 1977.
- Goldberg, David E. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley, 1989.
- MacLennan, Bruce J. "Evolution of Communication in a Population of Simple Machines." Technical Report CS-90-104, Computer Science Department, University of Tennessee, Knoxville, 1990.
- MacLennan, Bruce J. "Synthetic Ethology: An Approach to the Study of Communication." Technical Report CS-90-104, Computer Science Department, University of Tennessee, Knoxville, 1990.
- Rumelhart, D. E., G. E. Hinton, and R. J. Williams. Learning Internal Representations by Error Propagation." In *Parallel Distributed Processing*, edited by D. Rumelhart and J. McClelland, Volume 1. Cambridge, MA: MIT Press/Bradford Books, 1986.