

## Chapter X

# Environmental Variability and the Emergence of Meaning: Simulational Studies Across Imitation, Genetic Algorithms, and Neural Networks

Patrick Grim, State University of New York at Stony Brook, USA

Trina Kokalis, State University of New York at Stony Brook, USA

### Abstract

---

*A crucial question for artificial cognition systems is what meaning is and how it arises. In pursuit of that question, this paper extends earlier work in which we show the emergence of simple signaling in biologically inspired models using arrays of locally interactive agents. Communities of "communicators" develop in an environment of wandering food sources and predators using any of a variety of mechanisms: imitation of successful neighbors, localized genetic algorithms and partial neural net training on successful neighbors. Here we focus on environmental variability, comparing results for environments with (a) constant resources, (b) random resources, and (c) cycles of*

*"boom and bust." In both simple and complex models across all three mechanisms of strategy change, the emergence of communication is strongly favored by cycles of "boom and bust." These results are particularly intriguing given the importance of environmental variability in fields as diverse as psychology, ecology and cultural anthropology.*

## Introduction

Meaning is crucial to cognitive systems. It can be expected to be as crucial for artificial cognitive systems as it is for the ones we find occurring naturally around us, or indeed for the cognitive systems that we ourselves are. A crucial question for artificial cognition, then, is what meaning is and how it arises.

This paper is a development of earlier work in which we study the emergence of simple signaling in simulations involving communities of interacting individuals. Crucial to the model is an environment of wandering food sources and predators; our agents are "embodied" in an artificial environment and subject to its spatial and temporal contingencies. Crucial to the model is also the fact that it is not a single individual but a community of potentially interacting individuals that are embedded in such an environment. Our individuals develop coordinated behavioral strategies in which they make and respond to "sounds" in their immediate neighborhoods. Crucial to variations of the model explored here are different updating mechanisms of strategy change, all of which key to the behavior of most successful neighbors.

What our earlier work has shown, using any of various updating mechanisms in such a model, is the consistent emergence of communities of communicators using simple patterns of signaling. In an environment in which food sources and predators wander in a random walk, communities of individuals emerge that make a particular sound on successfully feeding, and respond to that same sound from neighbors by positioning themselves to feed. They make a different sound when hit by a predator, and respond to *that* sound from immediate neighbors by "hiding." Our models are biologically inspired in emphasizing strategy changes across a community of individuals embodied in a common environment. What consistently emerges are coordinated strategies of behavior that look a lot like simple signaling, and thus offer at least one clue to one kind of meaning.

What we introduce in this paper is a further characteristic of environments: variability. Our essential question is what role environmental variability — and environmental variability of what type — may play in the emergence of simple communication. Our inspiration comes from the role that environmental variability seems to play in a range of apparently disparate phenomena, from species diversity to individual learning.

In behavioral psychology, environmental variability has long been established as an important factor in operant conditioning. Intermittent schedules of reinforcement prove far more effective than constant reinforcement; variable-ratio schedules of reinforcement generally produce the highest number of responses per time period, establishing behavior most resistant to extinction (Reynolds, 1975; Honig & Staddon, 1977). "A

pigeon may peck the key 50 to 100 times after reinforcement has been cut off if it previously was on a schedule of continuous reinforcement. After some types of intermittent reinforcement, the bird will peck from 4,000 to 10,000 times before responding extinguishes" (Nye, 1992, p. 31).

In ecology and evolution, rates of environmental fluctuation have been proposed as a major factor in inter-species dynamics. A number of different mechanisms have been proposed linking environmental fluctuation to increased species diversity (Hutchinson, 1961; Harris, 1986; Huston, 1979; Hubbell & Foster, 1986; Chesson & Huntly, 1989, 1997). It has recently been proposed that Pleistocene climatic fluctuations are responsible for the evolution of larger brained mammals in general and higher primates in particular, with suggested links to social learning (Potts, 1996; Opdyke, 1995; Odling-Smee, Laland, & Feldman, 2000; Boyd & Richerson, 1985, 1989, 2000).

In cultural anthropology, variable environments appear to play a major role in the transition from foraging cultures to incipient agriculture. In a comprehensive computer model for archaeological data from the Guilá Naquitz cave site, R. G. Reynolds characterizes climate in terms of wet and dry years. Wet years show a wider range of food-acquisition behaviors with more new strategies, while dry years show a concentration on competitive and efficient strategies, with more critical pressure on strategy choice. Reynolds explains the role that environmental variability may play in the emergence of agriculture:

*The selective pressure placed on the group can vary unpredictably, and it is this variation that may be an important factor in determining the rate of change within the system. If, for example, the group was exposed only to a sequence of dry years that constantly put selective pressure on the group, the wet-year strategies that introduce most of the variation into the system would disappear or never be used. The resource scheduling system as a whole would be extremely conservative... On the other hand, an overbalance of wet years would introduce a good deal of variability into the system, but the group would seldom have an opportunity to test the worth of these adaptations in a more strenuous environment. (Reynolds, 1986, p. 499)<sup>1</sup>*

What our results here show is that environmental variability of a very specific sort plays a positive role in the simulational emergence of communication as well. Within spatialized models of self-serving individuals in an environment of wandering food items and predators, where each individual can hear and react to arbitrary sounds from immediate neighbors, a sine-wave variable environment with cycles of "boom and bust" promotes the development of communication. This effect appears regardless of important differences in how strategies are updated: Variable environments promote communication whether strategy change is by imitation of most successful neighbor, by genetic algorithm recombination with strategies of locally successful neighbors or by neural net training on successful neighbors.

Our previous work relied on biologically-inspired factors of environmental embodiment, emphasis on a community of individuals and development of coordinated behavior over time. What the results outlined in this paper indicate is that environmental variability of

a particular kind may be a further factor of importance in the development of meaning crucial to cognition systems.

More tentatively, we also want to offer a more speculative suggestion. It is tempting to think that the appeals to environmental variability across disciplines may have something in common. Perhaps there is some central mechanism of variability and selection which, in different forms, is responsible for aspects of individual learning, of species diversity in ecological communities and of the development of cultures.<sup>2</sup> That speculative suggestion is one that we will not pursue further here. It is an additional point in its favor, however, that environmental variability of a particular kind turns out to be of importance even in simulations of the emergence of simple patterns of meaning.

## The Basic Model

We work throughout with a 64×64 two-dimensional cellular automata array of 4,096 individuals, each of which follows a particular strategy. Initially these are chosen at random from our sample space of strategies (Figure 1). All action and reproduction are local within this spatialized array: Individuals interact only with their eight immediate neighbors, the cells that touch them at each side and at their corners. The array as a whole forms a torus, “wrapping around” so that individuals on the bottom edge have neighbors at the top edge and those at the left have neighbors on the right.


Individuals in the array alter their behavior in terms of what is happening immediately around them, but they do not move. In our simplest models, what move are food sources, which migrate in a random walk across the array. In our more complicated models we use both wandering food sources and wandering predators.

If a food source lands on an individual with its mouth open, that individual “feeds” and gains points. Individuals feed from food sources, but the sources are not consumed and don’t disappear. Like a cloud of plankton or a school of fish, perhaps, they continue their random walk across the array. In more complicated models, we include wandering predators and an appropriately more complicated repertoire of behaviors; individuals can open their mouths, “hide” or coast in neutral. An individual is “hurt” by losing a point if a predator lands on it when it isn’t hiding. In “neutral,” an individual fails to gain points from food sources but is still hurt if hit by a predator.

On any given round, an individual’s strategy may dictate that it opens its mouth or does not, where mouth-opening carries a particular cost in energy. In our more complex models, the strategy also dictates whether the individual hides or not, where “hiding” carries an energy cost as well. In all models, individual cells are capable of making sounds heard by their immediate neighbors. Sound-making, like mouth-opening and hiding, exacts an energy cost, but sounds come without any hard-wired significance: Nothing is built into the model in order to make a particular sound take on a particular meaning, or indeed to take on any meaning at all.

For even these simple individuals in this simple environment, there are behavioral strategies that seem to qualify as elementary forms of signaling or communication.

Figure 1. Initially randomized 64x64 array of 16 strategies



Strategies for different cells are shown in different colors, with open mouths shown as black central squares and migrating food sources as small white dots.


Imagine a spatially contiguous community of individuals that share the following strategy:

- They make a sound when they are successfully fed.
- They react to hearing a sound from their neighbors by opening their mouths.

When an individual in such a community feeds, it makes a sound. Its immediate neighbors, which share the same strategy, open their mouths in response. Since the food source continues its random walk, it will then fall on an open mouth on the next round. *That* individual, feeding successfully, will in turn make a sound and *its* neighbors will open their mouths in response. The result, in a community sharing such a strategy, is a chain reaction in which the food source is successfully exploited on each round (Figure 2). We term individuals with such a strategy "Communicators."

In our more complex models the environment contains both food sources and predators, and individuals can open their mouths, hide or coast in neutral on each round. In these

Figure 2. Migration of a single food source in a random walk across a hypothetical array of Communicators



In the left frame, a food source dot lands on an open mouth, indicated by gray shading. That central individual makes a sound \* heard by its immediate neighbors, which in the second frame open their mouths in response. One of them feeds successfully, making a sound heard by its immediate neighbors, which are shown opening their mouths in the third frame. The result is a community of Communicators is a chain reaction of efficient feeding.

models our agents have two sounds at their disposal, and a more complex form of communication is possible. Our “Perfect Communicators” are those that follow the following strategy:

- They make sound 1 when they are fed.
- They react to hearing sound 1 by opening their mouths.
- They make sound 2 when they are hurt.
- They react to sound 2 by hiding.

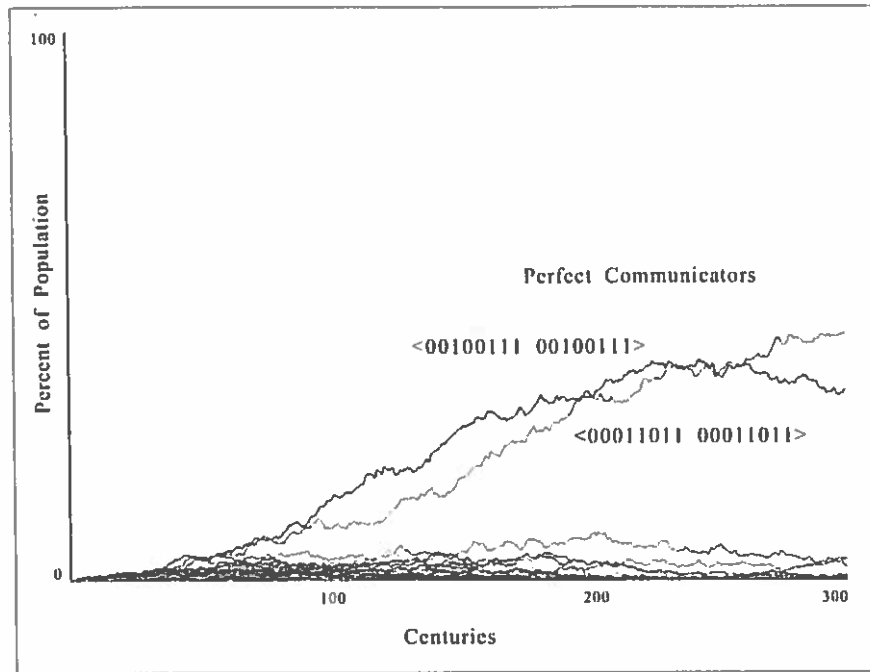
In previous work we have shown that these simple forms of communication can emerge from initially randomized arrays using any of several mechanisms for strategy change. In our earliest studies we used strategy change by simple imitation. At each “generation” — each 100 rounds of gain and loss from food capture and predation — each cell surveyed its immediate neighbors in order to see if any had garnered a higher score. If so, it changed to the strategy of its most successful neighbor (Grim, Kokalis, Tafti, & Kilb, 2000). In later studies we used strategy change by local genetic algorithm. Here the strategy of a less successful cell was replaced with a hybrid formed from its strategy and that of its most successful neighbor (Grim, Kokalis, Tafti, & Kilb, 2001). Most recently, we have instantiated strategies in the weights of simple neural nets, and have used strategy change by partial training on the behavior of more successful neighbors (Grim, St. Denis, & Kokalis 2002). Using any of these mechanisms in a wide variety of environments, we have been able to show that communities of Communicators will emerge and grow.<sup>3</sup> Figure 3, for example, shows a typical emergence of two forms of Communicators in an array of randomized neural nets over 300 generations. One of these Perfect Communicators uses sound 1 for food and sound 2 for predators; the other uses sound 2 for food and sound 1 for predators. Figure 4 plots the same results in terms of percentages of particular strategies within the population as a whole (Grim, Kokalis, Alai-Tafti, Kilb, & St. Denis, 2004).

*Figure 3. Emergence of two dialects of Perfect Communicators, shown in solid black and white, in a randomized array of simple neural nets with partial training on successful neighbors*



*Initially each of the 4,096 individuals in the array is randomly assigned one of seven discrete weights between -3.5 and + 3.5 for each of its 12 weights and biases. Cells are coded for different behaviors using background and dot color, with no visual representation for food sources, predators or reactions of opening mouths or hiding. Generations 1, 10, 50, 100, 200 and 300 are shown.*

Figure 4. Emergence of communication in a randomized array of simple neural nets with partial training on successful neighbors



Percentages of population graphed over 300 generations.

In all of our earlier studies, however, we considered sample spaces of strategies that came most easily with the strategy updating mechanism at issue. As a result we ended up with data on imitation runs, localized genetic algorithms and neural nets for quite different sample spaces of behaviors in a way that made it impossible to compare them side by side. In the current work we have deliberately structured our coded strategies and neural nets so as to produce the same range of possible strategies, allowing for a more direct comparison.<sup>4</sup>

In previous studies, we also used a constant environment of food sources and predators: Although each of these migrates in a random walk across the array, the total number of food sources and predators remains constant from generation to generation. Here, we focus instead on the role of a *variable* environment. Is change in the environment a factor of importance in the emergence of communication? Does the pattern of change matter, and does it perhaps have a different impact when different mechanisms of strategy change — imitation, localized genetic algorithms and spatialized neural nets — are in play?

The results that follow indicate that a variable environment does indeed have a major impact on the emergence of communication. The pattern of variability is crucial: Sine-

Copyright © 2007, Idea Group Inc. Copying or distributing in print or electronic forms without written permission of Idea Group Inc. is prohibited.

lid black  
ining on

ete weights  
t behaviors  
r reactions

ut written



wave variation in resources, with waves of boom and bust, has an effect that entirely random changes in resources with the same mean do not. In our studies that effect holds regardless of whether the basic mechanism of strategy change at issue is one of imitation, genetic recombination or neural net learning, and holds in similar ways in both simpler and more complex environments.

## The Simple Studies

---

Our initial studies use an environment with wandering food sources but without predators. The behavioral repertoire of our individuals is similarly limited: They can open their mouths or not, and can make a single sound heard by their immediate neighbors or remain silent. Individuals know only whether they have been successfully fed — whether a food source has landed on them when their mouth was open — and whether they or an immediate neighbor has made a sound on the previous round. Mouth opening carries an energy cost of .95 points, with an energy cost of .05 points for sounding.<sup>5</sup>

We code the behavior of these simple individuals in terms of four-tuples  $\langle f, \sim f, s, \sim s \rangle$ . Variable  $f$  dictates whether an individual makes a sound or not when it is fed,  $\sim f$  whether it makes a sound when it is not fed,  $s$  dictates whether it opens its mouth when it hears a sound from itself or an immediate neighbor and  $\sim s$  whether it opens its mouth when it hears no such sound.

This gives us only sixteen possible strategies, coded in binary as follows:

$\langle 0,0,0,0 \rangle$	Never makes a sound, never opens its mouth
$\langle 0,0,0,1 \rangle$	Never makes a sound, opens its mouth only when it hears no sound
$\langle 0,0,1,0 \rangle$	Never makes a sound, opens its mouth only when it hears a sound
$\langle 0,0,1,1 \rangle$	Never makes a sound, mouth always open
$\langle 0,1,0,0 \rangle$	Makes a sound when not fed, never opens its mouth
$\langle 0,1,0,1 \rangle$	Makes a sound when not fed, opens its mouth only when it hears no sound
$\langle 0,1,1,0 \rangle$	Makes a sound when not fed, opens its mouth only when it hears a sound
$\langle 0,1,1,1 \rangle$	Makes a sound when not fed, mouth always open
$\langle 1,0,0,0 \rangle$	Makes a sound when fed, never opens its mouth
$\langle 1,0,0,1 \rangle$	Makes a sound when fed, opens its mouth only when it hears no sound
$\langle 1,0,1,0 \rangle$	Makes a sound when fed, opens its mouth only when it hears a sound
$\langle 1,0,1,1 \rangle$	Makes a sound when fed, mouth always open

$\langle 1,1,0,0 \rangle$	Always sounds, never opens its mouth
$\langle 1,1,0,1 \rangle$	Always sounds, opens its mouth only when it hears no sound
$\langle 1,1,1,0 \rangle$	Always sounds, opens its mouth only when it hears a sound
$\langle 1,1,1,1 \rangle$	Always sounds, mouth always open

Those cells that carry strategy  $\langle 1,0,1,0 \rangle$  are our "Communicators." They make a sound when fed, and open their mouths when they hear a sound. A hypothetical community of Communicators will therefore behave as illustrated in Figure 2. We should also note that we use "imperfect" worlds throughout. All cells follow their programmed strategies subject to a 5% measure of error. In 5% of those cases in which the strategy specifies mouth opening, the mouth fails to open; in 5% of cases where it is specified as opening, it stays shut. Nowak and Sigmund (Nowak & Sigmund, 1990, 1992) have argued that a measure of stochastic "noise" makes for a more realistic model of cooperation. In previous work we have outlined its importance for the emergence of communication as well (Grim, Kokalis, Tafti, & Kilb, 2000).

This sample space of behaviors remains the same across our simple studies. These behaviors are instantiated in different ways in different cases, however — as coded behaviors or as operating neural nets. This allows us to compare different mechanisms for strategy change side by side: strategy change by imitation, by localized genetic algorithm and by localized training of neural nets.


In one series of runs our individuals carry behaviors coded as series of binary digits and follow an imitation algorithm for strategy change. After 100 rounds of food gathering, point gain and energy loss, each cell surveys its immediate neighbors and sees if any has garnered a higher score. If so, it adopts the strategy of its highest-scoring neighbor in place of its own.

In a second series of runs, we use the same coding for behaviors but employ a localized genetic algorithm for strategy change. After 100 rounds, each cell surveys its immediate neighbors to see if any has garnered a higher score. If not, it retains its current strategy. If it has a more successful neighbor, however, that cell's strategy is replaced with a genetic algorithm hybrid formed from its current strategy and that of its most successful neighbor. We use two-point crossover, choosing one of the offspring at random to replace the parent (Figure 5).

Genetic algorithms are usually applied globally to a population, breeding from only a small number of those strategies that perform most successfully on some uniform fitness function. Ours in contrast is a *localized* genetic algorithm. All genetic recombination is local: Cells with locally successful neighbors change their strategies to local hybrid recombinations. Unlike global genetic algorithms, localized genetic algorithms seem promising here not merely as a sampling device — the sample space in our simple studies is only 16 strategies — but as a means of facilitating strategy similarity and thus the possibility of behavioral coordination between neighbors.

In a third series of runs we generate the same sample space of behaviors using very simple neural nets (Figure 6). On each round, an individual has either heard a sound from one of its immediate neighbors or it has not, coded as a bipolar input of +1 or -1. It has also


Figure 5. Genetic recombination of simple strategies



successfully fed on that round or not, again coded as a bipolar +1 or -1. Our neural structure involves just two weights and two biases, each of which carries a value between -3.5 and +3.5, "chunked" at one-unit intervals. A bipolar input at "hear sound," for example, is multiplied by weight  $w_1$ . To that product is added the bias value, which might equivalently be thought of as a weight on an input of +1. The two values are summed at the output node. If the result is greater than a threshold of 0, the output is treated as +1; the individual opens its mouth, for example. If it is less than or equal to 0, the output is treated as -1 and the individual keeps its mouth closed.

These simple nets generate the same range of behaviors as their coded counterparts, but they allow for a different mechanism of strategy change. For our neural nets, strategy change is by partial training on successful neighbors. After 100 rounds, each cell surveys its immediate neighbors to see if any has garnered a higher score. If so, it does a partial training on the behavior of its most successful neighbor. A single training consists of a random pair of inputs for both the "trainee" cell and its more successful neighbor. If the two nets give the same output, no change is made in the trainee's weights. If the outputs are not the same, the trainee's weights are nudged a single unit toward what would have given its neighbor's response on that run. Biases are shifted in the same way. With bipolar coding and within the limits of our value scale, using "target" for the

Figure 6. Simple neural nets



Simple neural nets with bipolar inputs of +1 or -1 at "hear sound" and "fed," multiplied by weights  $w_1$  and  $w_2$  "chunked" at one-unit intervals between -3.5 and +3.5. Biases carry similar weights. If the total at the output node  $> 0$ , the output is treated as +1 and the individual opens its mouth, for example. If the total output is  $\leq 0$ , the output is treated as -1 and the individual keeps its mouth closed.

neighbor's output, we can calculate the delta rule as simply  $w_{new} = w_{old} + (\text{target} \times \text{input})$  and  $\text{bias}_{new} = \text{bias}_{old} + \text{target}$ .

Our aim was to compare these three different mechanisms — imitation, localized genetic algorithm and neural net training — across a range of different environments. In one series of runs, we used environments constant with regard to food resources: Exactly 50 food sources migrated across the array each generation. In a second case, we used randomly variable environments. Here the average was again 50 food sources each generation, but the specific number of food sources at any given generation was a random number between 0 and 100. In a third form of variability, we assigned food sources by sampling at regular intervals along a sine wave oscillating between 0 and 100. This again gave us an average of 50 food sources each generation, but in waves of increasing and decreasing resources each generation.

The core question was whether these differences in environmental variability would make a difference in the emergence of communication. We also wanted to know whether any such difference would depend on a particular mechanism of strategy change.

## Environmental Variability and the Emergence of Communication

---


### Constant Environment

---

Our constant environments contained exactly 50 food items each time. We used the gain allotted for successful feeding as an independent variable: Tests were run with gains from 1 to 140 points for each successful feeding. What we wanted to plot was what strategy an array would evolve to for particular gains — to Communicators or otherwise — and in what number of generations. We could then compare results across mechanisms of imitation, localized genetic algorithms and neural nets.

Figure 7 shows results across different gains for the imitation algorithm, in which strategies simply switch to those of their most successful neighbors. Starting from the left of the graph, we sample gains of 1 through 9, shown in narrow bars. From that point we switch to sampling gains at 10-point spreads, with the wider bars representing gains from 10 through 140. Runs are to 1,500 generations; the height of each bar indicates how many generations were required for fixation on a single strategy across the entire array. Should no single strategy occupy the entire array by 1,500 generations, the bar tops out. For bars below 1,500, then, color indicates the strategy in total occupation of the array. For bars that reach 1,500, color indicates the dominant strategy across the array at that point. Dominant strategies at each gain are also indicated by their codes in the list at the left. In a constant environment, the initial winner at most gains below 9 is the null strategy  $\langle 0,0,0,0 \rangle$ , which neither sounds nor opens its mouth. Strategy  $\langle 0,0,1,0 \rangle$ , a “free rider” which responds to sound by opening its mouth but reciprocates with no sound in return, makes a brief appearance at gains of 6 and 8 points. Starting at a gain of 7, however, and

Figure 7. Results across different gains for strategy replacement by imitation



Strategy replacement by imitation of most successful neighbor in a constant environment of 50 food sources, showing runs up to 1,500 generations with gains from 1 to 140 points for each successful feeding. Columns below 1,500 generations show total conquest by strategy; those at 1,500 show dominant strategy at that point. All Eat conquers or proves dominant at gains of 100 and above.

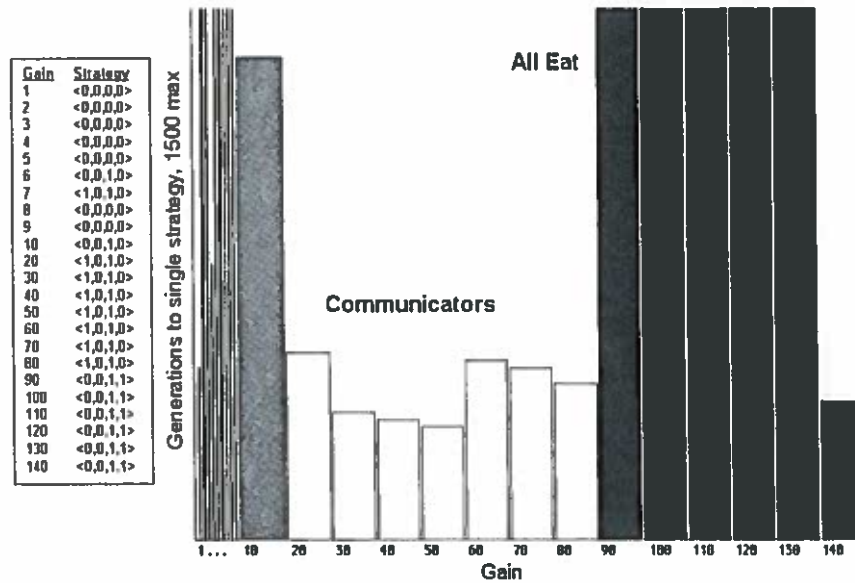
then from gains of 9 through 90, it is our Communicators  $\langle 1,0,1,0 \rangle$  to which the array converges.<sup>6</sup>

It is immediately obvious from the graph, and somewhat surprising, how large the window for communication is. Communicators dominate the array from the point at which each successful feeding is worth 10 points to the point at which it is worth 9 times as much. But it is also clear that communication has an upper terminus: Above a gain of 100 points it is a strategy of All Eat that proves dominant. Beyond this point all arrays are dominated by  $\langle 0,0,1,1 \rangle$ , which doesn't bother to communicate at all: It sits with a constant open mouth, ready to catch passing food sources, but never making a sound.

Strategy change by localized genetic algorithm in the same constant environment and for the same set of gains gives us the results shown in Figure 8. For genetic algorithm recombination, as for imitation, there is a clearly circumscribed window for communication. Here the window is somewhat smaller, extending only from gains of 20 to gains of 80. Below that dominance goes to the null strategy  $\langle 0,0,0,0 \rangle$  or the "free rider"  $\langle 0,0,1,0 \rangle$ . At gains of 90 and above dominance again goes to the incommunicative All Eat.

Though the overall pattern is similar, the genetic algorithm runs take several times longer. This is perhaps not too surprising. In an imitation model, a successful strategy is imitated

Figure 8. Strategy replacement by localized genetic algorithm




Strategy replacement in an environment by localized genetic algorithm combination with most successful neighbor in a constant environment of 50 food sources, showing runs up to 1,500 generations with gains from 1 to 140 points for each successful feeding. Columns below 1,500 generations show total conquest by strategy; those at 1,500 show dominant strategy at that point. All Eat conquers or proves dominant at gains of 90 and above.

immediately by its neighbor. In a localized genetic algorithm, complete cloning to a successful neighbor may take a significant number of hybrid recombinations. Where a large number of possible strategies is at issue, one advantage to genetic algorithm recombination is that it can explore a larger portion of the strategy space: Strategies will be produced which don't exist in either parent. Where we are dealing with a sample space of only 16 strategies, that general advantage doesn't show to effect.

The window for communication in an array of neural nets with constant environment is closely comparable to arrays updating strategy by imitation and localized genetic algorithm. Here the null strategy <0,0,0,0> dominates at gains of 8 and below for each successful feed. Communicators dominant for gains between 9 and 80, but All Eat proves dominant at gains of 90 and above.

Despite greater complexity and partial training, our neural nets reach fixation in this simple environment at about the same rate as mere imitators where Communicators dominate, and much faster where dominance is by All Eat.<sup>7</sup> It should be noted, however, that speed is comparable only when measured in numbers of generations; actual computer time for each run was significantly longer in the case of neural nets.

Figure 9. Strategy replacement by partial neural net training



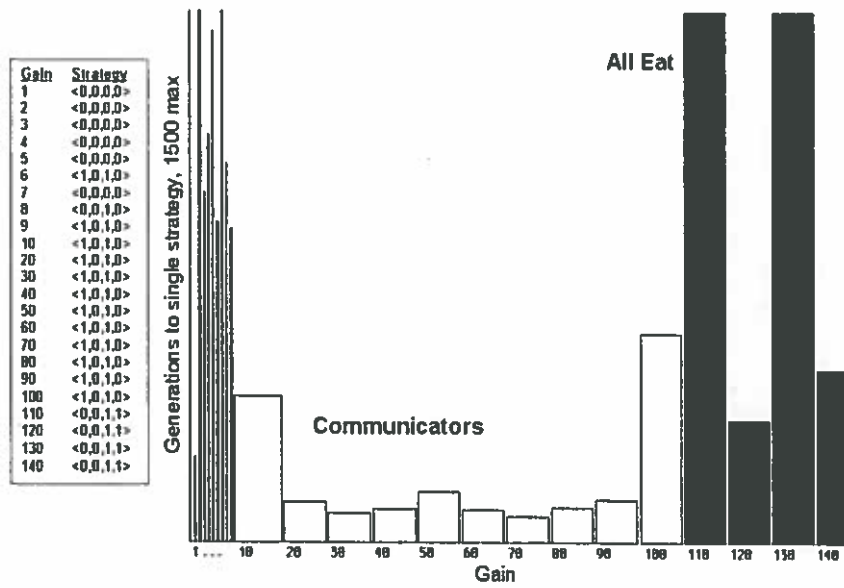
Strategy replacement by partial neural net training on most successful neighbor in a constant environment of 50 food sources, showing runs up to 1,500 generations with gains from 1 to 140 points for each successful feeding. Columns below 1,500 generations show total conquest by strategy; those at 1,500 show dominant strategy at that point. All Eat conquers or proves dominant at gains of 90 and above.

## Random Environment


What if, instead of a constant environment of 50 food sources, we use a changing environment? In a second series of studies we assigned a random number of food sources between 0 and 100 each generation. The average number of food sources remained at 50, but the particular number of food sources on any generation might be anywhere between 0 and 100. The amount of gain allotted for successful feeding was again our independent variable: Tests were run with gains for each successful feeding from 1 to 140 points for each successful feeding. Figure 10 shows results in a random environment for strategy change by imitation, localized genetic algorithm and neural nets.

With any of our mechanisms of strategy change, it turns out, results in a randomized environment show at most a slight increase in the upper limit for Communicators. With constant food sources, we found an upper limit of 90, 80 and 80 as the gains at which Communicators proved dominant for strategy change by imitation, localized genetic algorithm and neural nets respectively. With a randomized number of food sources that lifts slightly to upper limits of 100, 90 and 90 for the three cases. In both constant and randomized environments, however, the window for Communicators closes at a gain of 90 or 100 and All Eat proves dominant from that point on.

Figure 10. Strategy replacement in an environment of randomized food sources



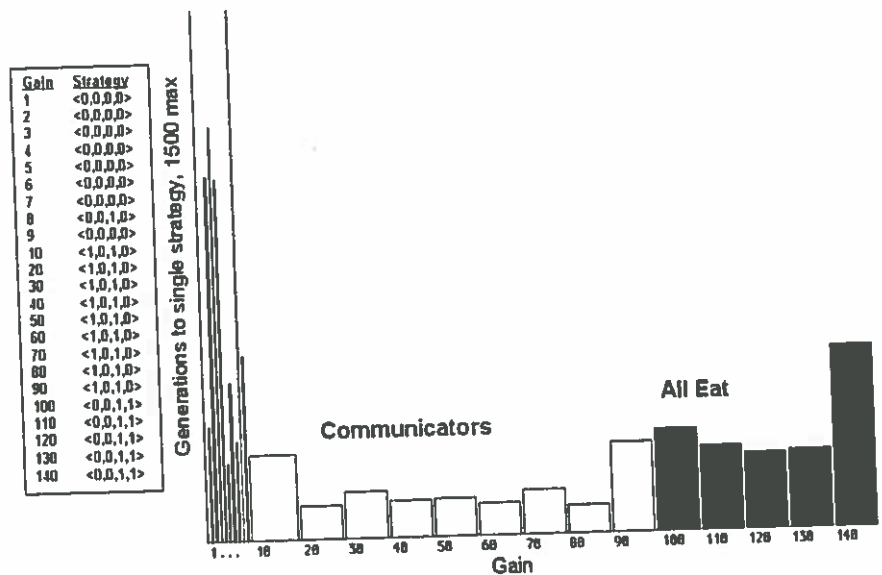
(a) Imitation in a random environment



(b) Localized genetic algorithm in a random environment



Figure 10. continued



(c) Neural nets in a random environment

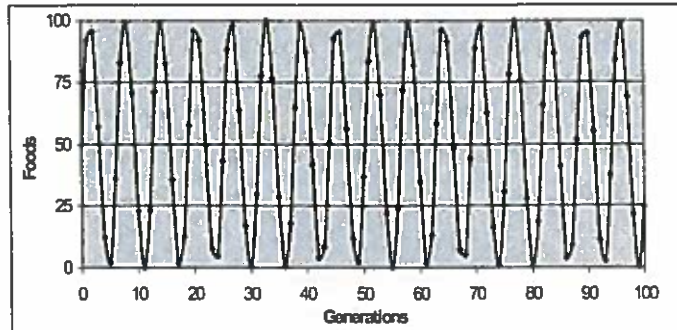
Strategy replacement in an environment of randomized food sources with a mean of 50 by (a) imitation of most successful neighbor, (b) localized genetic algorithm combination with most successful neighbor and (c) partial neural net training on most successful neighbor. Runs are shown up to 1,500 generations with gains from 1 to 140 points for each successful feeding. Columns below 1,500 generations show total conquest by strategy; those at 1,500 show dominant strategy at that point. All Eat conquers or proves dominant at gains of 100 or 110 and above in each case.

## Sine-Wave Variable Environment

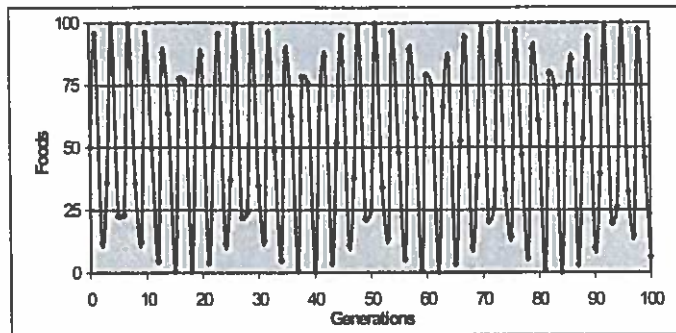
An environment with a random number of food sources produces much the same effects as one with a constant number of food sources. But what if we use an environment which, though variable, shows greater regularity? What if there is a cycle of "boom and bust," for example — will this make a difference in the emergence of communication?

The decision to test environments with "boom and bust" cycles still leaves a great deal of latitude, since there may be very different patterns qualifying as "boom and bust." We conceived of different patterns in terms of different intervals marked out on a regular sine wave oscillating between 0 and 100. With values of that wave taken at intervals of 1, we get one pattern of numbers for our food sources. With values taken at intervals of 2, we get a different series.

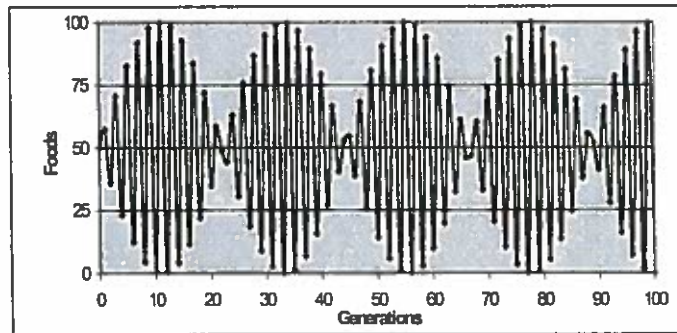
Figure 11. Different patterns of variation in the number of food sources in the environment over time



(a) Sin+1 variability



(b) Sin+2 variability



(c) Sin+3 variability

Each data point in each graph, representing number of food sources in the array at that generation, is given by  $[\sin(x) + 1] * 50$  food sources for a series  $x_0, x_1, \dots, x_n$ . In each series  $x_{n+1} = x_n + c$  for a specific  $c$ . In (a),  $x_{n+1} = x_n + 1$ . In (b),  $x_{n+1} = x_n + 2$ . In (c),  $x_{n+1} = x_n + 3$ .

138 148

of 50 by (a)  
n with most  
is are shown  
columns below  
ategy at that  
h case.

ame effects  
nent which,  
n and bust,"  
on?


a great deal  
id bust." We  
regular sine  
vals of 1, we  
vals of 2, we

thout written

More formally, our number of food sources each time was dictated by a regular series  $x_0, x_1, \dots, x_n$ , where each  $x_n$  adds some constant  $c$  to its predecessor  $x_{n-1}$ . We take our number of food sources as  $[\sin(x) + 1] * 50$  for this series, giving us some number of food sources between 0 and 100 for each generation. The distance  $c$  we choose between elements of our series then dictates the particular form of "boom and bust."

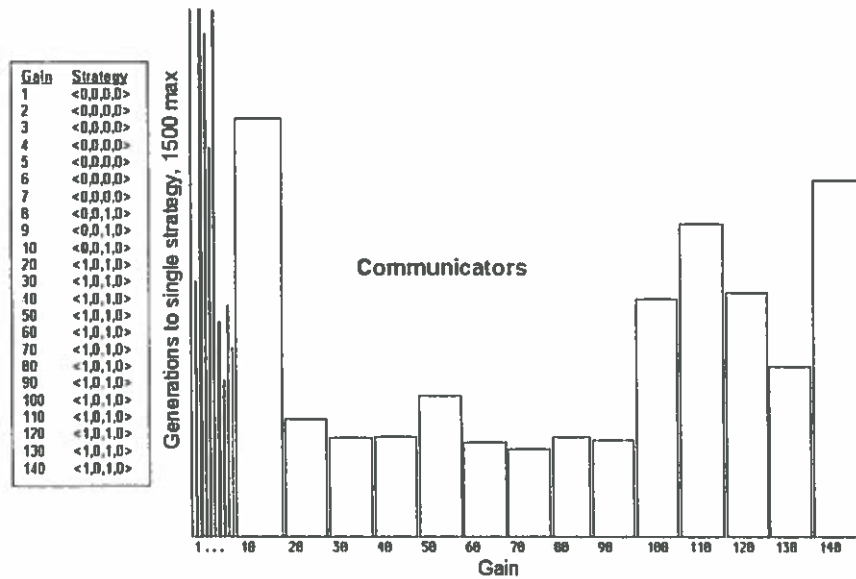
Figure 11 shows different patterns of variation. Each data point in the first graph is  $[\sin(x) + 1] * 50$  for a series  $x_0, x_1, \dots, x_n$  that increments by units of one. In the first graph, in other words,  $x_{n+1} = x_n + 1$ . In the second graph our series increments by units of two:  $x_{n+1} = x_n + 2$ . In the third graph we increment by units of three:  $x_{n+1} = x_n + 3$ . Although we average 50 food sources in each case, the differences in boom and bust patterns are clear. A sine value for a series which changes by increments of +1 gives a fairly gradual change between boom and bust, with one or two intermediate points between a high and a low. Sine value for a series which changes by increments of two show a more polarized boom and bust, with a midpoint present in only half of the transitions from top to bottom. Sine values for a series which changes by increments of three swing instantaneously from boom to bust without any midpoints, though the interval distance between boom and bust progressively enlarges and narrows in cycles over time. Because of their more dramatic shifts, we used the second and third patterns of variation as our samples, referring to these patterns simply as  $\sin+2$  and  $\sin+3$ .

Figure 12. Triumph of Communicators at all gains above 10 in an environment of food sources between 0 and 100 varying in the pattern of  $\sin+2$




(a) Imitation in a  $\sin+2$  environment

Figure 12. continued



(b) Localized genetic algorithms in a Sin+2 environment



(c) Neural nets in a Sin+2 environment


Strategy replacement is by (a) imitation of most successful neighbor, (b) localized genetic algorithm combination with most successful neighbor and (c) partial neural nets training on most successful neighbor.

What impact does a sine-wave variable environment have on the emergence of communication? Figure 12 shows emergence of communication in an environment changing on the pattern of  $\sin+2$  for each of our three mechanisms of strategy change. The surprising result is that a variable environment allows conquest by Communicators all the way up. Unlike constant and random environments, increased gains in a variable environment on the pattern of  $\sin+2$  do not favor All Eat at any point within the scope of the graph. We have sampled larger gains, beyond the scope of the graph; up to gains of 500 and 1,000 it is still the Communicators that succeed.

The result is sensitive to patterns of variability — the more gradual changes of a  $\sin+1$  pattern do not show as dramatic a result. The effect on communication is by no means confined to the pattern of  $\sin+2$ , however. Similar results for resources following the pattern of  $\sin+3$  are shown in Figure 13.


Across all of our modes of strategy change, sine-wave variable environments of these patterns show a dramatic widening of the window of gain values in which Communicators appear and flourish. Although the average number of food sources remains the same as in our constant and randomly variable environments, cycles of “boom and bust” strongly favor the emergence of communication.<sup>8</sup>

Figure 13. Triumph of Communicators at all gains above 10 in an environment of food sources between 0 and 100 varying in the pattern of  $\sin+3$

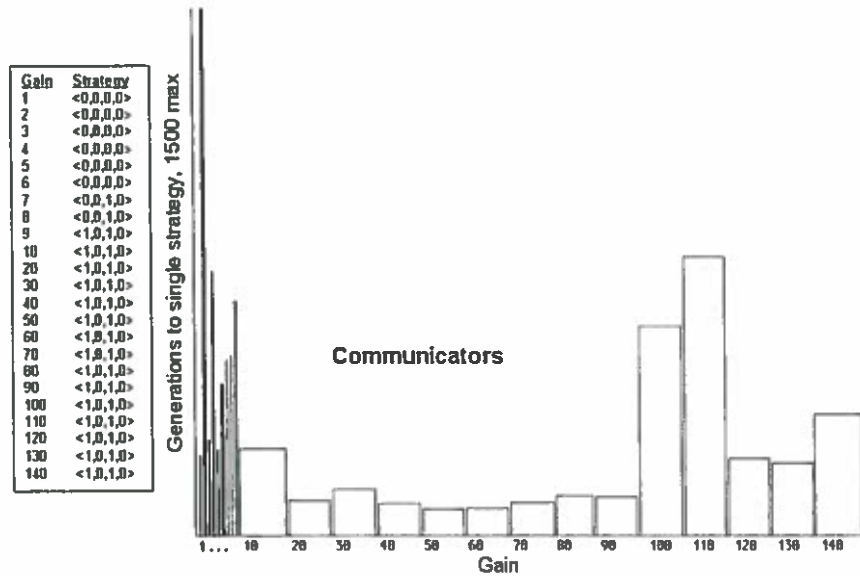


(a) Imitation in a  $\sin+3$  environment

Figure 13. continued



(b) Localized genetic algorithms in a Sin+3 environment



(c) Neural nets in a Sin+3 environment

Strategy replacement is by (a) imitation of most successful neighbor, (b) localized genetic algorithm combination with most successful neighbor and (c) and partial neural nets training on most successful neighbor.

## The Complex Studies

---

We extend the model into a range of more complex studies in which environments contain not only food sources but predators. Here we use 75 food sources and 150 predators, each of which wanders in a random walk across the array.<sup>9</sup>

In this second series of studies, our individuals have a larger behavioral repertoire. On any given round they can open their mouths, hide or coast in neutral. An individual is awarded a particular gain if it "feeds" — if its mouth is open when a food source lands on it, but also loses a point when "harmed" — when it is hit by a predator and is not hiding. Mouth-opening and hiding each carry an energy cost of .05 points.<sup>10</sup> Coasting in neutral carries no energy cost, but when in neutral an individual can neither benefit by feeding nor avoid predation by hiding. Here our individuals also have two arbitrary sounds at their disposal rather than one. Sound-making, like mouth-opening and hiding, carries an energy cost of .05 points.

We again compare three mechanisms of strategy change: imitation, localized genetic algorithm and partial training of neural nets. In the case of imitation and genetic algorithm, we encode our strategies as ternary six-tuples  $\langle \sigma, f, h, s1, s2, s\sigma \rangle$ . The first three variables specify what sound an individual makes when neither fed nor hurt  $\sigma$  (no sound, sound 1 or sound 2), what sound it makes when fed  $f$  (the same three options) and what sound it makes when hurt  $h$  (the same three options). The second three variables specify what action an individual takes when it hears sound 1 (coast in neutral, open its mouth or hide), when it hears sound 2 and when it hears no sound  $s\sigma$  (the same three options in each case).


For simplicity, we allow our individuals to respond only to being fed, hurt or neither; there is no provision for responding to being both fed or hurt. We also have our individuals respond to only one sound or neither, again with no provision for both. If an individual is both fed and hurt, or hears both sounds 1 and 2, we randomize which input it responds to. We structure individuals so that they can make only one sound at a time and can engage in only one action — mouth-opening or hiding. With these restrictions we can keep our specifications to six-tuples  $\langle \sigma, f, h, s1, s2, s\sigma \rangle$ , with a sample space of 729 possible strategies. Here, as before, we should also mention that an element of "noise" is built in: An individual will open its mouth in a random 5% of cases even if its code specifies otherwise, and will similarly hide in a random 5% of cases.

Where strategy change is by imitation, each individual looks around to see if a neighbor has acquired a higher score after 100 rounds of point gain and loss. If so, it adopts the strategy of its most successful neighbor. Should there be more than one neighbor with equal higher scores, the strategy of one is chosen randomly.

Using a localized genetic algorithm, here as before, we hybridize the strategy of an individual with that of its highest scoring neighbor, should any neighbor prove more successful. We use two-point crossover on our ternary six-tuples, choosing one of the offspring at random (Figure 14).

Our neural nets are structured as the simple perceptrons shown in Figure 15. This two-lobe structure for communication has been re-invented or re-discovered repeatedly in the history of the literature. Since De Saussure (1916), many have noted an intrinsic distinction between (1) making sounds or sending signals, and (2) responding to sounds

Figure 14. Genetic recombination applied to strategies of six variables




or signals received. It seems natural to embody that distinction in the neural architecture of the individuals modeled.<sup>11</sup> Here we use eight weights with four biases, each of which carries a value between -3.5 and +3.5, “chunked” at one-unit intervals. We use a bipolar coding for inputs, so that “hear sound 1” takes a value of +1 if the individual hears sound 1 from an immediate neighbor on the previous round, for example, and takes a value of -1 if it does not. Each input is multiplied by the weight shown on arrows from it, and the weighted inputs are then summed with the weight of the bias at the output node. Were we operating these nets “pure,” their inputs at “hear sound 1” and “hear sound 2,” for example, could both be +1. Their outputs at “open mouth” and “hide” could also both be +1. In order to allow comparison with our coded behaviors in imitation and genetic algorithm runs, we constrain inputs and outputs. In operation, we “cripple” our nets. We program our simulation so that only one sound or neither can be fed into a net as input on any given round, with a similar constraint on “fed” and “hurt” inputs. Where there are two sounds in the environment, or a cell is both fed or hurt, one input is chosen at random much as in the imitation and genetic algorithm models. Even with these input limitations, however, it would be possible for nets with particular patterns of weights to give a positive output for both “make sound 1” and “make sound 2,” for example, or for both “open mouth” and “hide.” Here we change our threshold for outputs. Should neither output give a weighted sum >0, both are treated as outputting -1, with the result that no sound is made, for example. Should only one output give a sum >0, that output is treated as the sole output of +1. Should both outputs give a sum >0, we treat that with the greater sum as +1 and the other as -1; when both are >0 and equal, we pick one at random. Here, as in the other cases, we also build in an element of “noise”; in a random 5% of cases individuals will open their mouths regardless of weights and inputs, hiding in a random 5% of cases as well.

We must admit that these constraints on our neural nets are in some way “unnatural”; without those constraints they would show a greater range of behaviors in a greater variety of input situations. We also speculate that these constraints contributed to the slowness with which our nets operated and the slowness with which large arrays of such nets evolved.<sup>12</sup>

In our neural net runs, as in the others, our individuals total their points over the course of 100 generations. If any of their eight immediate neighbors has garnered a higher score,



Figure 15. Neural nets in the more complex model, using eight weights and four biases for each individual



Bipolar inputs of +1 or -1 are taken at "hear sound 1," "hear sound 2," "fed" and "hurt," and are multiplied by the weights on arrows from them. Biases carry similar weights. "Open mouth," "hide," "make sound 1" and "make sound 2" are the output nodes. If the total at an output node  $> 0$ , the output is treated as +1 and the individual "opens its mouth" or "makes sound 2," for example. If the total output is  $\leq 0$ , the output is treated as -1 and the individual keeps its mouth closed or does not make sound 2.

the cell's highest-scoring neighbor is taken as its "target." For a set of four random inputs, with no provision against duplication, the cell then compares its outputs with those of its target. At any point at which outputs differ, each weight and bias in the "trainee" is nudged a single unit toward what would have given its neighbor's response on that run. With bipolar coding, we can calculate changes in weights and biases as  $w_{\text{new}} = w_{\text{old}} + (\text{target} \times \text{input})$  and  $\text{bias}_{\text{new}} = \text{bias}_{\text{old}} + \text{target}$ . Any such training will clearly be partial. Only four sets of inputs are sampled, rather than the full 16 possible, and indeed the same set may be sampled repeatedly. The delta rule is applied using each set of inputs only once, moreover, leaving no guarantee that each weight will be shifted enough to match the target behavior. The idea of partial training is deliberately built into the model in order to allow behavioral strategies to emerge that might not have existed in either trainee or target.

In each form of strategy change we are now dealing with the same 729 behaviors, representable in ternary notation as six-tuples  $\langle o, f, h, s1, s2, s\emptyset \rangle$ . Of these, there are only two that qualify as "Perfect Communicators."  $\langle 0, 1, 2, 1, 2, 0 \rangle$  makes no sound when neither fed nor hurt, makes sound 1 when fed and makes sound 2 when hurt. It correspondingly responds to hearing sound 1 by opening its mouth, and responds to hearing sound 2 by hiding.  $\langle 0, 2, 1, 2, 1, 0 \rangle$ , on the other hand, uses sound 2 when fed and sound 1 when hurt, responding to sound 2 by opening its mouth and to sound 1 by hiding.

In this more complicated environment, however, with a significantly wider range of possible behaviors, a number of broader categories of behavior turned out to be important:

- *All Eat*: A total of 27 of our coded behavioral strategies fell in this category: Any with a string of three 1s on the right-hand side of the code. These strategies vary in the sounds they produce when they are fed or hurt. But they don't vary in the

action they take in response to any sound pattern; regardless of whether they hear sound 1, sound 2 or neither, they sit with open mouths. Our All Eat strategies therefore include:

```

000111 100111 200111
001111 101111 201111
002111 102111 202111
010111 110111 210111
011111 111111 211111
012111 112111 212111
020111 120111 220111
021111 121111 221111
022111 122111 222111
    
```

- *All Hide:* These 27 strategies differ in that their final three digits are 2s rather than 1s:

```

000222 100222 200222
001222 101222 201222
002222 102222 202222
010222 110222 210222
011222 111222 211222
012222 112222 212222
020222 120222 220222
021222 121222 221222
022222 122222 222222
    
```

- *Eat Default Communicators:* These eight strategies have some signal operating as a predator warning — sending sound 1 when hurt and responding to sound 1 by hiding, for example — but only against a background default to “open mouth” in the absence of any signals. They therefore have a “1” in the furthest column to the right. This category includes communicators that use and respond to symbols for both food and predators, but only against an “open mouth” default.

```

001201 001211
001221 002021
002121 002221
012121 021211
    
```

- *Hide Default Communicators:* Here it is "hiding" that is the default, with at least communication signals for the presence of food.

010102 010112  
010122 012122  
020012 010112  
020212 021202

- *Food Communicators:* These six strategies communicate perfectly with each other, including a no-action "neutral" in the absence of signals, but communicate only about the presence of food. Their code may specify a response to a signal that is "idle" in the sense that they don't send it, but there is no signal sent in the absence of feeding and no default action in the absence of signals.

010100 010110  
010120 020010  
020110 020210

- *Predator Communicators:* These six strategies communicate perfectly, but only about the presence of predators.

001200 001210  
001220 002020  
002120 002220

- *Perfect Communicators:* Two strategies send and receive signals about both food and predators, coasting in neutral in the absence of any signals.

012120 021210

In our simple studies, we used the gain allotted for successful feeding as the independent variable with which to measure the "window" for communication in different environments. We use the same measure in our more complex studies, keeping "loss" for predation constant at 1 point but varying the "gain" for successful feeding from 0 to 150. In a first series of runs, across all modes of strategy change, we used a constant environment of precisely 75 wandering food sources and 150 wandering predators. In a second series, we used an environment with the same averages for food sources and predators, but with the number of food sources picked as a random number between 1 and 150 each generation and the number of predators picked at random between 1 and

300. In a third series, we used an environment with the same average for each but with coordinated sine-wave variation for numbers of food sources and predators. Here, as before, the question was whether these differences in environment would make a difference in the emergence of communication.


### Constant Environment

Gain for feeding is again our independent variable, but is sampled at larger intervals than before: here at 1, 5, 10, 25, 50, and 150 points for each successful feeding. Because of the complexities of a larger sample space of strategies, we concentrated not on particular strategies but on which category of strategies proved dominant at each gain. Results for strategy change by simple imitation in a constant environment are shown in Figure 16.

With a gain of 1, Perfect Communicators occupy the entire array of 4,096 cells in 375 generations. For gains of 10, 25, 50, 100, and 150, however, the array goes to fixation with all cells playing All Eat. The fact that our bar goes to the top of the graph for a gain of 5 indicates that no single strategy category occupied the entire array by our limit of 1,500 generations. In this case the color of the bar indicates only the category dominant in the array at that point: In this case, for example, Eat Default Communicators occupied 3,842 cells of the 4,096. All Eat occupied 254 cells. A graph of the particular dynamics in that case is shown in Figure 17.

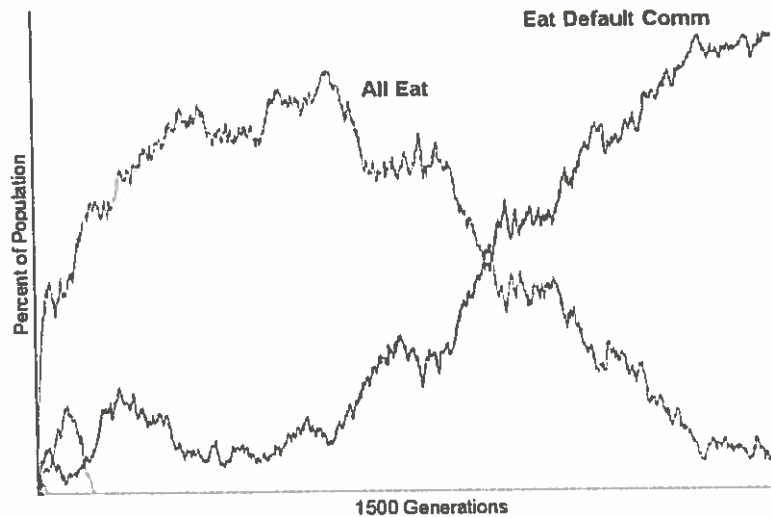
For strategy change by imitation, then, a constant environment favors All Eat for all gains over 5. All Eat also proves dominant when strategy change is by localized genetic algorithm (Figure 18).

Figure 16. Dominance by All Eat with strategy change by imitation



Dominance by All Eat with strategy change by imitation of most successful neighbor in a constant environment of food sources and predators. Generations to conquest by a single strategy category shown for different gains, 1,500 generations maximum.

Figure 17. Competition between All Eat and Eat Default Communicators for a gain of 5; 1,500 generations of strategy change by imitation in a constant environment




What of neural nets? Throughout our trials, arrays of the more complex but constrained neural nets took significantly longer to come to fixation. For none of our gain settings did we arrive at a single strategy category within 1,500 generations. Figure 19 is thus remarkably dull, showing merely that All Eat was the dominant category at 1,500 generations for all gains greater than 1. Here, another way of graphing results proves more informative. Figure 20 shows relative populations for different categories established by generation 1,500 at each gain. Bars don't extend to the end, it should be noted, because of an assortment of strategies that don't fall into any of our listed categories.

Despite the slowness of our nets, this too appears to be a relative victory for All Eat. At 1,500 generations, and for gains of 5 and above, All Eat is the only one of our behavioral categories that has established itself in a significant proportion of the population. Perfect Communicators and Food Communicators appear only for very small gains.

### Random Environment

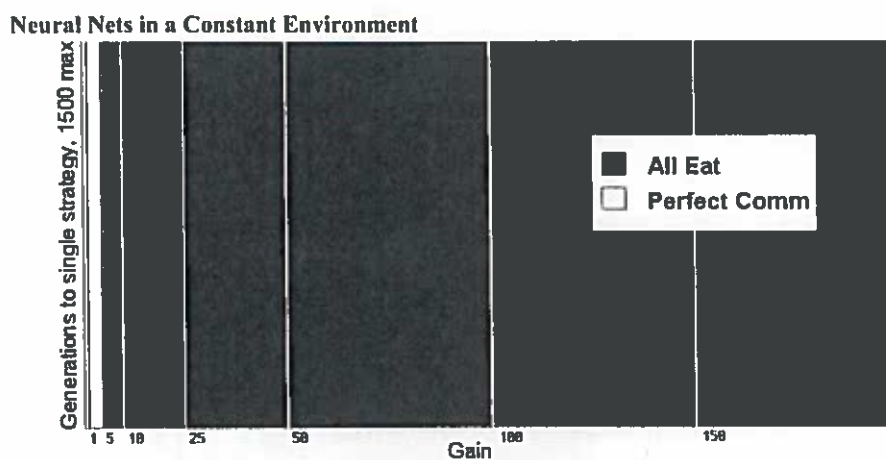
For the random runs, we set food sources at a random number between 1 and 150 each time, with predators at a random number between 1 and 300. These averaged to 75 food sources and 150 predators over the course of a run, but of course varied unpredictably from generation to generation. Results for each mode of strategy change are shown in Figure 21.

Figure 18. Dominance by All Eat with strategy change by localized genetic algorithm



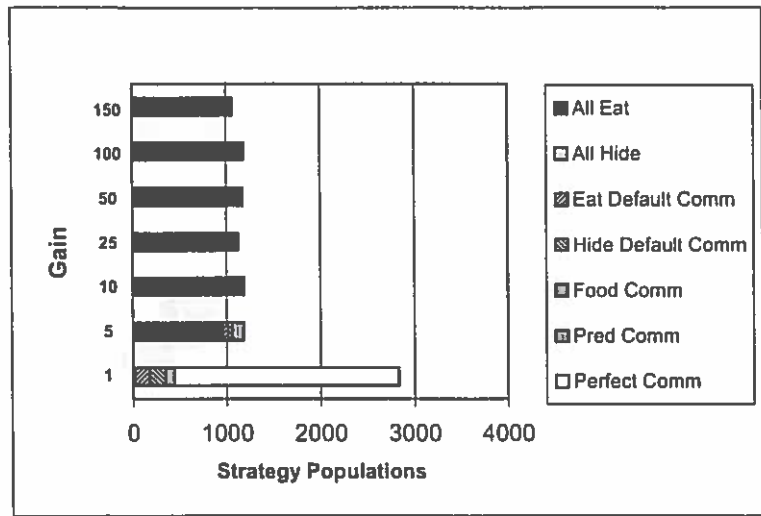
Dominance by All Eat with strategy change by localized genetic algorithm combination with most successful neighbor in a constant environment of food sources and predators.

Figure 19. Dominance by All Eat with strategy change by localized neural net training




Dominance by All Eat with strategy change by localized neural net training on most successful neighbor in a constant environment of food sources and predators. Colors of bars at 1,500 indicate strategy dominance but not necessarily full conquest.

Figure 20. Domination by All Eat: Populations of strategy categories




Domination by All Eat: Populations of strategy categories at different gains for partial neural net training on most successful neighbor in constant environment of food sources and predators.

Figure 21. Categories dominant in an environment of randomized numbers of food courses and predators




(a) Imitation in a random environment



(b) Localized genetic algorithms in a random environment

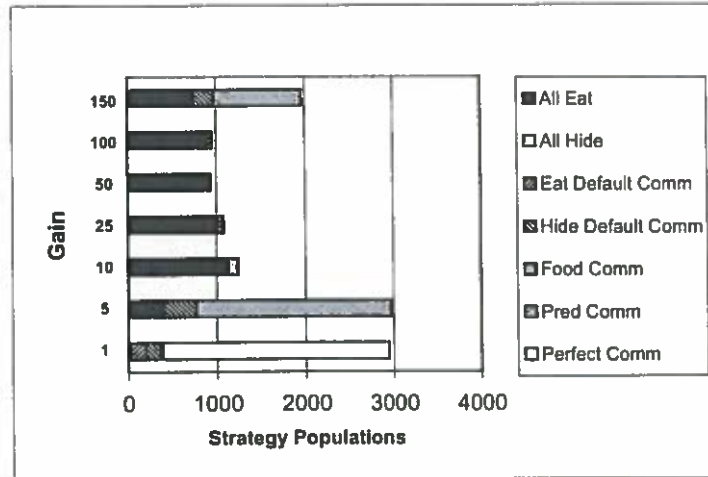
Figure 21. continued



(c) Neural nets in a random environment

Categories dominant in an environment of randomized numbers of food sources and predators. Strategy replacement by (a) imitation of most successful neighbor, (b) localized genetic algorithm combination with most successful neighbor and (c) partial neural nets training on most successful neighbor. Colors of bars at 1,500 indicate strategy dominance but not necessarily full conquest.

Figure 22. Populations of strategies in major categories at different gains for partial neural net training in an environment with random numbers of food sources and predators





In the imitation runs, Perfect Communicators take over with a gain of 1, but at greater gains the general winners are Eat Default Communicators. This is a family of strategies that qualify as "Communicators" because of reciprocal signaling regarding predators. Like All Eat, however, the winner at a gain of 50, they maintain an open mouth in the absence of any signal.

In the genetic algorithm runs, Food Communicators prove dominant with gains of 1 and 5, though they don't completely occupy the array within 1,500 generations. Food Communicators resemble Perfect Communicators in that they send no signal when neither fed nor hurt, and take no action when no signal is received. But their communication is limited to the presence of food. Beyond a gain of 5 it is All Eat that takes over.

Populations of strategies in major categories at different gains for partial neural net training in an environment with random numbers of food sources and predators.

In no case did our neural nets reach fixation on a single strategy category within 1,500 generations. Here again results for neural nets are clearer in Figure 22. Perfect Communicators are dominant at a gain of 1, with Food Communicators at a gain of 5. Beyond that, the dominant category is All Eat, with the peculiar exception of close competition between Food Communicators and All Eat at a gain of 150.

In an environment of randomized numbers of food sources and predators, there do seem to be important differences in the outcomes of our models depending on whether we use imitation, localized genetic algorithm or partial training of neural nets. In the case of imitation, in particular, Eat Default Communicators make a significant showing that doesn't appear with either of our other modes of strategy change. In neural nets and genetic algorithms, Food Communicators also make a more significant showing than they did in a constant environment. All Eat, however, is still the most significant category for gain levels above 5 in both neural nets and genetic algorithms.

### **Sine-Wave Variable Environment**

---

In our simple studies, communication was strongly favored in sine-wave variable environments. Will the same hold for the more complex model?

With one major qualification, the answer is "yes": Here again, cycles of "boom and bust" strongly favor the emergence of communication. The major qualification has to do with the precise form of communication that is favored. Our more complex models include both food sources and predators, with the possibility of communicative strategies regarding either or both. Numbers of both food sources and predators vary together in cycles of "boom and bust."<sup>13</sup> The variable we use to test the "window" for communication, however — level of gain for successful feeding — is clearly relevant only to communication regarding food.

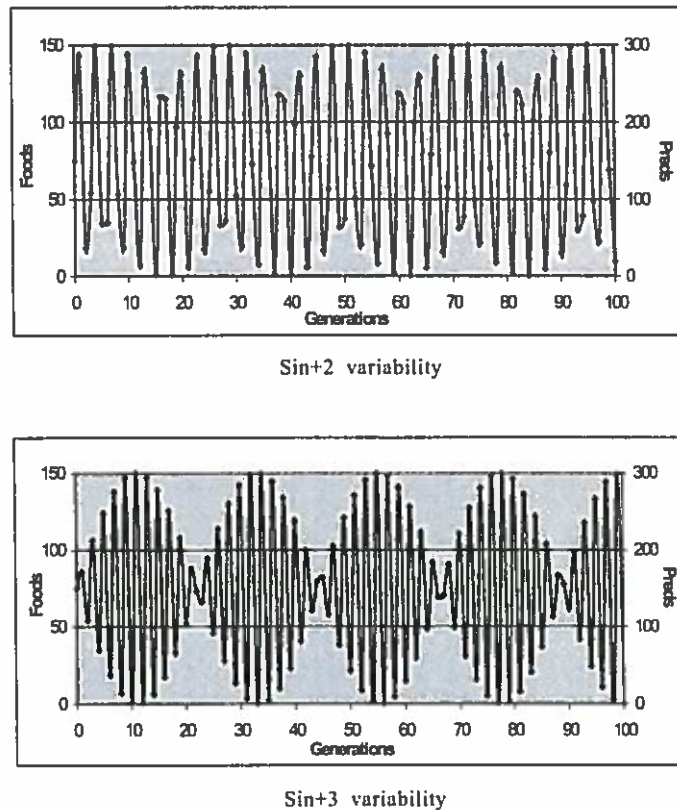
Communication regarding food, it turns out, *is* strongly favored in sine-wave variable environments: Food Communicators dominate, unthreatened by All Eat, at very high levels of gain. Because loss for predation remains constant at one penalty point across our runs, on the other hand, communication regarding predation does not appear. With up to 150 points to be gained by successful feeding, the one point that might be lost by

predation proves insufficient to produce either Pred Communicators or fully Perfect Communicators.

Here, as before, we explored two patterns of sine-wave variability, shown in Figure 23. In this case, data points represent both numbers of food sources between 1 and 150 and predators between 0 and 300. In the sin+2 graph, our food sources are plotted as  $[\sin(x) + 1] * 75$  for a series  $x_0, x_1, \dots, x_n$  which increments by units of two; predators are plotted as  $[\sin(x) + 1] * 150$  for the same series. In the second series, our increments are by units of three. Although we compiled full data for both forms of variation, the final results prove nearly identical. We therefore exhibit only the results for the sin+2 series as representative of both.<sup>14</sup>

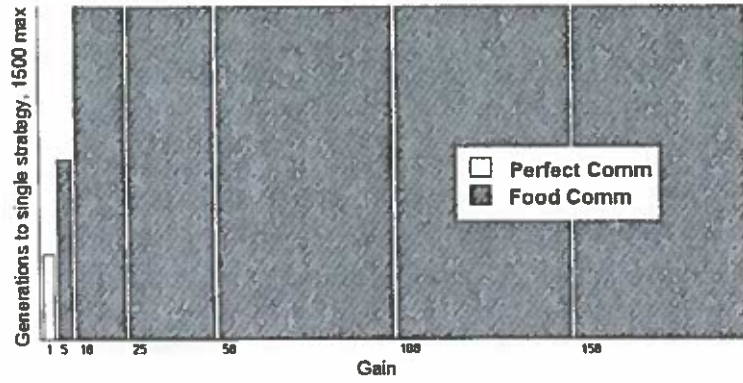
Because almost all runs failed to reach fixation by 1,500 generations, results appear in a less informative form in Figure 24. What these graphs do make clear is the complete dominance by Food Communicators for all modes of strategy change and all gains greater than one. Proportions of populations in each case are shown in Figure 25.

Figure 23. Patterns of variation for food sources and predators

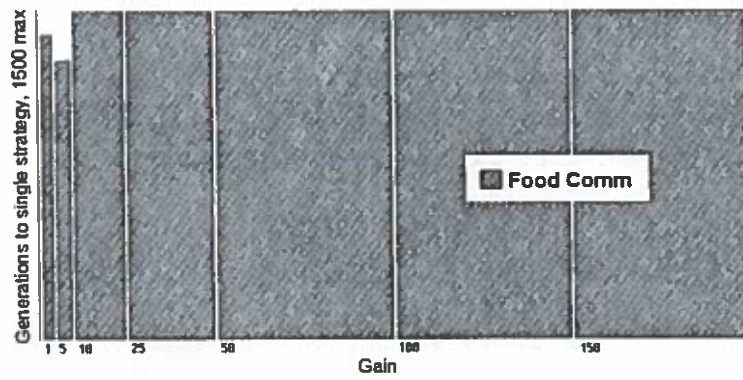


$[\sin(x) + 1] * 75$  for food sources and  $[\sin(x) + 1] * 150$  for predators with different incremental series  $x_0, x_1, \dots, x_n$ . In the top series,  $x_{n+1} = x_n + 2$ . In the second series,  $x_{n+1} = x_n + 3$ .


Figure 24. Full dominance by Food Communicators across all modes of strategy change in a sine-wave variable environment



(a) Imitation in a Sin+2 environment




(b) Localized genetic algorithm in a Sin+2 environment



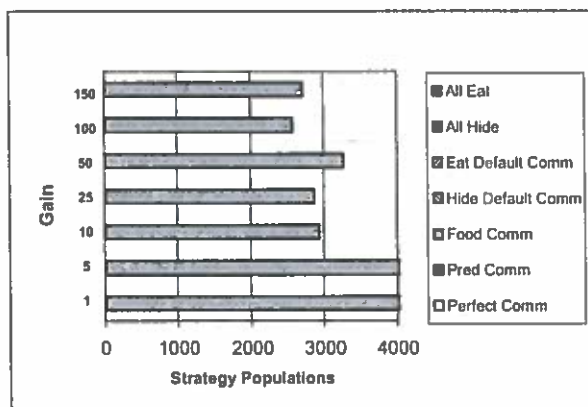
(c) Neural nets in a Sin+2 environment

Results shown are for sin+2. Sin+3 results are nearly identical.

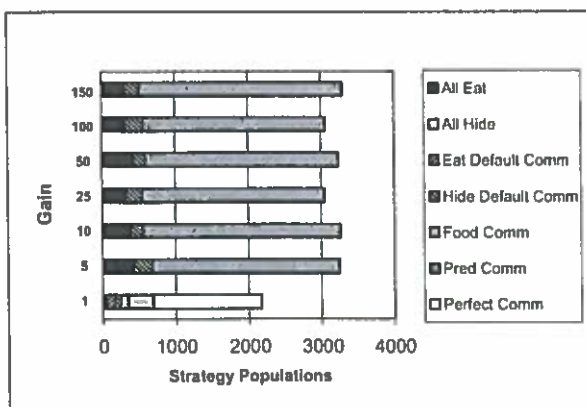
Figure 25. Full dominance by Food Communicators across all modes of strategy change in a  $\sin+2$  variable environment



(a) Imitation in a Sin+2 environment



(b) Localized algorithms in a Sin+2 environment



(c) Neural nets in a Sin+2 environment

Populations of strategies at 1,500 generations. Sin+3 results are nearly identical.

We take these results as clear evidence that sine-wave variable environments of at least the "boom and bust" form of  $\sin+2$  and  $\sin+3$  strongly favor the emergence of communication. In constant and random environments dominance at gains above 5 goes to either All Eat or Eat Default Communicators. Food communicators play a subsidiary role. With a variable environment of the forms explored here, on the other hand, dominance is by Food communicators for gains from 5 to 150.

Our more complex simulations therefore underscore the lesson of the simpler models above. Across a range of modes of strategy change, sine-wave variable environments can play a major role in the emergence of communication.

## Conclusion

---

In earlier studies we found that communities of Communicators can emerge from an initially randomized array of strategies in an environment of wandering food sources and predators. Communication can emerge, moreover, using any of three different mechanisms of strategy change: imitation of successful neighbors, combination by localized genetic algorithm with most successful neighbors and partial neural net training on the behavior of most successful neighbors (Grim, Kokalis, Tafti, & Kilb, 2000, 2001; Grim, St. Denis, & Kokalis, 2002).

In our earlier studies, however, we used environments with a constant number of wandering food sources and predators. Here our attempt has been to expand those studies to questions of environmental variation: Is communication about resources more favored in an environment in which the level of resources are variable than in an environment in which resources are constant?

For an environment with randomly variable resources, the answer is "no." In both our simpler and more complex simulations, random variation in resources showed much the same effect as constant resources with the same average. In an environment with sine-wave variable resources, on the other hand — an environment of "boom and bust" resource cycles — the answer is clearly "yes." It is thus not merely variability but the particular pattern of variability that is of importance; communicative strategies are much more strongly favored in sine-wave variable environments. That effect holds whether the mechanism of strategy change at issue is one of imitation, localized genetic algorithm or partial training on neural nets.

The advantage to communication in a variable environment, we speculate, is that the behavior of a community of Communicators can be environmentally sensitive. Unlike many of their competitors, communities of Communicators can take effective advantage of "boom" cycles and yet harbor their energy resources in times of "bust." A more thorough understanding of the mechanisms of environmental variability in this kind of simulation, however, as well as a wider exploration of different patterns of variation, will require further work.

Meaning is crucial to cognitive systems, whether artificial or natural. Our earlier work suggests that a key to understanding meaning may lie in understanding the development

of behavioral coordination in communities of agents embodied in an environment. The current results suggest that environmental variability may also play an important role in the emergence of meaning. In terms of natural cognition systems, these results offer hints of a deeper understanding of some of the mechanisms of meaning. They may also offer hints toward richer development of meaning in artificial cognition.

It is intriguing that environmental variability has been appealed to as an important explanatory factor in a range of different disciplines. In ecology, environmental fluctuation has been seen as playing an important role in species diversity (Hutchinson, 1961; Harris, 1986; Huston, 1979; Hubbell & Foster, 1986; Chesson & Huntly, 1989, 1997). In cultural anthropology, cycles of boom and bust have been linked to the growth of agriculture (Reynolds, 1986). Pleistocene climatic fluctuations have recently been proposed as instrumental in the evolution of larger brained mammals in general and higher primates in particular, with speculative links to social learning and culture (Potts, 1996; Opdyke, 1995; Odling-Smee, Laland, & Feldman, 2000; Boyd & Richerson, 1985, 1989, 2000). The study of particular patterns of environmental variability and their impact is perhaps most developed in decades of careful work on schedules of reinforcement and operant conditioning (Reynolds, 1975; Honig, & Staddon, 1977; Nye, 1992). It is tempting to speculate that these appeals to environmental variability across disciplines may have some central mechanism in common. We take it as a suggestive fact, worthy of further investigation, that environmental variability turns out to be important even in the simple simulational studies of communication and meaning we have outlined here.

## References

- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: Chicago University Press.
- Boyd, R., & Richerson, P. J. (1989). Social learning as an adaptation. *Lectures on mathematics in the Life Sciences*, 20, 1-26.
- Boyd, R., & Richerson, P. J. (2000). Built for speed: Pleistocene climate variation and the origin of human culture. In F. Tonneau & N. S. Thompson (Eds.), *Perspectives in ethology, Vol. 13: Evolution, culture, and behavior* (pp. 1-45). New York: Kluwer Academic/Plenum.
- Cangelosi, A., & Parisi, D. (1998). The emergence of a 'language' in an evolving population of neural networks. *Connection Science* 10, 83-97.
- Chesson, P. L., & Huntly, N. (1989). Short-term instabilities and long-term community dynamics. *Trends in Ecology and Evolution*, 4, 293-298.
- Chesson, P. L., & Huntly, N. (1997). The roles of harsh and fluctuating conditions in the dynamics of ecological communities. *American Naturalist*, 150, 519-553.
- De Saussure, F. (1916). *Cours de linguistique generale*. R. Harris (Trans. 1983). *Course in general linguistics*. London: Duckworth.
- Flannery, K. V. (1986). Adaptation, evolution, and archaeological phases: Some implications of Reynolds' simulation. In K. V. Flannery (Ed.), *Guilá naquitz: Archaic*

- foraging and early agriculture in Oaxaca, Mexico* (pp. 501-507). New York: Academic Press.
- Grim, P. (1995). Greater generosity in the spatialized Prisoner's Dilemma. *Journal of Theoretical Biology*, 173, 353-359.
- Grim, P. (1996). Spatialization and greater generosity in the stochastic Prisoner's Dilemma. *Biosystems*, 37, 3-17.
- Grim, P., Kokalis, T., Alai-Tafti, A., Kilb, N., & St. Denis, P. (2004). Making meaning happen. *Journal of Experimental and Theoretical Artificial Intelligence*, 16, 209-243.
- Grim, P., Kokalis, T., Tafti, A., & Kilb, N. (2000). Evolution of communication in perfect and imperfect worlds. *World Futures: The Journal of General Evolution*, 56, 179-197.
- Grim, P., Kokalis, T., Tafti, A., & Kilb, N. (2001). Evolution of communication with a spatialized genetic algorithm. *Evolution of Communication*, 3, 105-134.
- Grim, P., Mar, G., & St. Denis, P. (1998). *The philosophical computer*. Cambridge, MA: The MIT Press.
- Grim, P., St. Denis, P., & Kokalis, T. (2002). Learning to communicate: The emergence of signaling in spatialized arrays of neural nets. *Adaptive Behavior*, 10, 45-70.
- Harris, G. P. (1986). *Phytoplankton ecology, structure, function, and fluctuation*. London: Chapman & Hall.
- Honig, W. K., & Staddon, J. E. R. (Eds.). (1977). *Handbook of operant behavior*. Englewood Cliffs, NJ: Prentice Hall.
- Hubbell, S. P., & Foster, R. B. (1986). Biology, chance, history and the structure of tropical rainforest communities. In J. Diamond & T. J. Case (Eds.), *Community ecology* (pp. 314-329). New York: Harper and Row.
- Huston, M. (1979). A general hypothesis of species diversity. *American Naturalist*, 113, 81-101.
- Hutchinson, G. E. (1961). The paradox of the plankton. *American Naturalist*, 95, 145-159.
- MacLennan, B. J. (1991). Synthetic ethology: An approach to the study of communication. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II, SFI Studies in the Sciences of Complexity* (Vol. X, pp. 631-655). Redwood City, CA: Addison-Wesley.
- Nowak, M. A., Plotkin, J. B., & Jansen, V. A. A. (2000). The evolution of syntactic communication. *Nature*, 404, 495-498.
- Nowak, M. A., Plotkin, J. B., & Krakauer, D. C. (1999). The evolutionary language game. *Journal of Theoretical Biology*, 200, 147-162.
- Nowak, M., & Sigmund, K. (1990). The evolution of stochastic strategies in the prisoner's dilemma. *Acta Applicandae Mathematicae*, 20, 247-265.
- Nowak, M., & Sigmund, K. (1992). Tit for tat in heterogeneous populations. *Nature*, 355, 250-252.
- Nye, R. D. (1992). *The legacy of B. F. Skinner: Concepts and perspectives, controversies and misunderstandings*. Pacific Grove, CA: Brooks/Cole.

- Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. (2000). Niche construction and gene-culture coevolution: An evolutionary basis for the human sciences. In F. Tonneau & N. S. Thompson (Eds.), *Perspectives in ethology: Vol. 13. Evolution, culture, and behavior* (pp. 89-111). New York: Kluwer Academic/Plenum.
- Oliphant, M., & Batali, J. (1997). Learning and the emergence of coordinated communication. *Center for Research on Language Newsletter*, 11(1).
- Opdyke, N. (1995). Mammalian migration and climate over the past seven million years. In E. S. Vrba, G. H. Denton, T. C. Partridge, & L. H. Burckle (Eds.), *Paleoclimate and evolution with emphasis on human origins* (pp. 8-23). New Haven, CT: Yale University Press.
- Potts, R. (1996). *Humanity's descent: The consequences of ecological instability*. New York: William Morrow.
- Reynolds, G. S. (1975). *A primer of operant conditioning*. Glenview, IL: Scott, Foresman & Company.
- Reynolds, R. G. (1986). An adaptive computer model for the evolution of plant collecting and early agriculture in the eastern valley of Oaxaca. In K. V. Flannery (Ed.), *Guilá naquitz: Archaic foraging and early agriculture in Oaxaca, Mexico* (pp. 439-500). New York: Academic Press.
- Seth, A. K. (1998). The evolution of complexity and the value of variability. In C. Adami, R. Belew, H. Kitano, & C. Taylor (Eds.), *Artificial Life VI: Proceedings of the Sixth International Conference on the Simulation and Synthesis of Living Systems*. Cambridge, MA: The MIT Press.
- Werner, G., & Dyer, M. (1991). Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II, SFI Studies in the Sciences of Complexity* (Vol. X, pp. 659-687). Redwood City, CA: Addison-Wesley.

## Endnotes

- <sup>1</sup> See also Flannery (1986).
- <sup>2</sup> Questions of environmental variability are applied to robotics and the evolution of complexity in Seth (1998).
- <sup>3</sup> For earlier work with a similar model regarding cooperation rather than communication (see Grim, 1995, 1996, and Grim, Mar, & St. Denis, 1998).
- <sup>4</sup> Although we haven't pursued it here, we take direct comparison of different computational structures to be of both practical and theoretical importance. In terms of practical implications, new program optimization algorithms are constantly being proposed, including new variations on some of those used here. It is to be expected that some of these will prove more successful in the optimization of some kinds of programs, under some conditions, while others prove more successful for other kinds of programs or under other conditions. Only rarely, however, has there



been any attempt to "bench test" optimization algorithms side by side. With regard to theoretical implications, it is clear that animal behavior can be shaped by different mechanisms, including, for example, genetics and individual learning. Will one form of mechanism prove optimal for a particular range of behaviors, for a particular batch of organisms or in a particular set of conditions, while another mechanism turns out to be optimal for other behaviors or organisms in other conditions? Here a number of theoretical questions seem ripe for simulational studies using direct comparisons between different mechanisms.

- <sup>5</sup> These energy costs were carried over from earlier studies that did not address environmental variability but in which they proved auspicious for the development of communication at low gains. Results across different costs are detailed in Grim, Kokalis, Tafti, and Kilb (2000, 2001), and Grim, St. Denis, and Kokalis (2002).
- <sup>6</sup> Because of limited computer resources and the extensive computing time required for some runs across ranges of gains, particularly those involving neural net training, graphs throughout represent results of single runs at each gain. In a more comprehensive study, it would clearly be desirable to average multiple runs at each point. Consistency of results across ranges of gains, however, offers similar assurance of reliability — the fact Communicators triumph consistently at gains from 10 through 90 in Figure 7, for example.
- <sup>7</sup> We do not have a tidy explanation for the quicker convergence to All Eat using neural nets.
- <sup>8</sup> A speculation as to why this result holds is offered in the conclusion.
- <sup>9</sup> The reason for using twice as many predators as food sources is detailed in Grim, Kokalis, Tafti, and Kilb (2000). A bit of reflection on the dynamics of feeding and predation built into the model shows an important difference between the two. In an array composed entirely of Communicators, as illustrated in Figure 2, a chain reaction can be expected in terms of food signals and successful feedings. The dynamics of a "hurt" alarm are very different. Among even Perfect Communicators, a cell signals an alarm only when it is hurt — that is, when a predator is on it and it isn't hiding. If successful, that "alarm" will alert a cell's neighbors to hide, and thus the predator will find no victim on the next round. Precisely because the predator then finds no victim, there will be no alarm sounded, and thus on the following round even a fellow "communicator" may be hit by a predator. Here one sees not a chain reaction of successful feeding on every round, but an alternating pattern of successful avoidance of predation every second round. An important difference between the dynamics of feeding and the dynamics of predation is thus built into the model. With equal numbers of food sources and predators, that difference in dynamics would strongly favor communication regarding food over communication regarding predators. One way to compensate for that difference is simply to proportion food sources and predators accordingly, as we have done here.
- <sup>10</sup> The .05 "tax" for mouth-opening is significantly less than the .95 tax in the simpler studies because there is an additional down-side to mouth-opening in this richer environment. When its mouth is open, an individual cannot be hiding and so is vulnerable to predators.

- <sup>11</sup> This maps precisely onto the distinction between “emissions” and “actions” in MacLennan (1991) and between “transmission behavior” and “reception behavior” in Oliphant and Batali (1997). These two functions are separated between two different sexes in Werner and Dyer (1991), and between two separate sets of connection weights in the neural nets of Cangelosi and Parisi (1998). Martin Nowak notes that an active matrix for signal-sending and a passive matrix for signal-reading can be treated as completely independent in Nowak, Plotkin, and Krakauer (1999) and in Nowak, Plotkin, and Jansen (2000).
- <sup>12</sup> Results below were often compiled using multiple computers running sub-sets of gains over a period of weeks and months.
- <sup>13</sup> The possibility of varying food sources and predators independently, and at different rates, remains an open question. The initial and limited explorations we made in this direction gave no clear or decisive results.
- <sup>14</sup> Full data for both series are of course available on request.