

Boom and Bust: Environmental Variability Favors the Emergence of Communication

Patrick Grim and Trina Kokalis

Group for Logic & Formal Semantics, Dept. of Philosophy, SUNY at Stony Brook
pgrim@notes.cc.sunysb.edu

Abstract

Environmental variability has been proposed as an important mechanism in behavioral psychology, in ecology and evolution, and in cultural anthropology. Here we demonstrate its importance in simulational studies as well. In earlier work we have shown the emergence of communication in a spatialized environment of wandering food sources and predators, using a variety of mechanisms for strategy change: imitation (Grim, Kokalis, Tafti & Kilb 2000), localized genetic algorithm (Grim, Kokalis, Tafti & Kilb 2001), and partial training of neural nets on the behavior of successful neighbors (Grim, St. Denis & Kokalis 2002). Here we focus on environmental variability, comparing results for all of these mechanisms in a range of different environments: (a) environments with constant resources, (b) environments with random resources around the same mean, and (c) sine-wave variable environments with cycles of 'boom and bust'. Communication, it turns out, is strongly favored by environmental variability on the pattern of 'boom and bust'.

Introduction

Sometimes the same idea appears persistently across a range of very different disciplines. Environmental variability seems to be such an idea.

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, with variable-ratio schedules producing the highest number of responses per time period and establishing behavior most resistant to extinction (Honig & Staddon 1966, Nye 1992).

In ecology and evolution, rates of environmental fluctuation have been proposed as a major factor in inter-species dynamics (Chesson and Huntly 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 (Opdyke 1995, Potts 1996, Boyd & Richerson 2000).

In cultural anthropology, variable environments appear to play a major role in the transition from foraging cultures to incipient agriculture (Reynolds 1986).

It is tempting to think that these appeals to environmental variability may have something in common. Perhaps there is some central mechanism of

variability and selection which, in different forms, is responsible for the way that individuals learn, the way that species evolve, and the way that cultures develop. The results we offer here indicate that an important role for environmental variation shows up even in simple computer simulations.

The Basic Model

We work throughout with an initially randomized 64 x 64 two-dimensional cellular automata array of 4,096 individuals carrying different behavioral strategies (Figure 1). All action and reproduction are local: individuals interact only with their eight immediate neighbors.




Figure 1 Randomized array of strategies

Our individuals alter their behavior in terms of what is happening immediately around them, but they do not move. In the simple models offered here, it is food sources that move, migrating in a random walk across the array. If a food source lands on an individual with its mouth open, that individual 'feeds' and gains points. Our food sources are not consumed, however; like a cloud of plankton or a school of fish, they continue their random walk, offering nourishment for the next individual down the line.

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. Its strategy also dictates whether or not it makes a sound on that round, heard by itself and its immediate neighbors. Sound-making also carries a cost in energy.

For even these simple individuals in this simple environment, there are forms of behavior that would seem to qualify as elementary forms of signaling or

communication. Imagine a 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. 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’.




Figure 2 Migration of a food source across an array of Communicators. Gray shading indicates an open mouth, * the range of a sound made.

In previous work, using both this model and a more complex variation incorporating predators, we have shown that simple forms of communication can emerge from initially randomized arrays using any of several mechanisms for strategy change. In our earliest studies we use strategy change by simple imitation. After 100 centuries of gain and loss, each cell surveys its immediate neighbors in order to see if any had garnered a higher score. If so, it adopts the strategy of its most successful neighbor (Grim, Kokalis, Tafti, and Kilb 2000). In later studies we use strategy change by local genetic algorithm. Here the strategy of a less successful cell is replaced with a hybrid, created by genetic algorithm from its strategy and that of its most successful neighbor (Grim, Kokalis, Tafti, and 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, and Kokalis 2002). Using any of these mechanisms, we have been able to show that communities of Communicators will emerge and grow. Figure 3 shows a typical emergence of two forms of Communicators—here with signals for both food and predators—in an array of randomized neural nets over 300 generations.

In this background research, however, we use an environment of constant resources: although our food sources migrate in a random walk across the array, the total number of food sources 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? The

results that follow indicate that a variable environment does indeed have a major impact on the emergence of communication, even in computer simulations as simple as those explored here. The pattern of variability, it turns out, is also crucial.

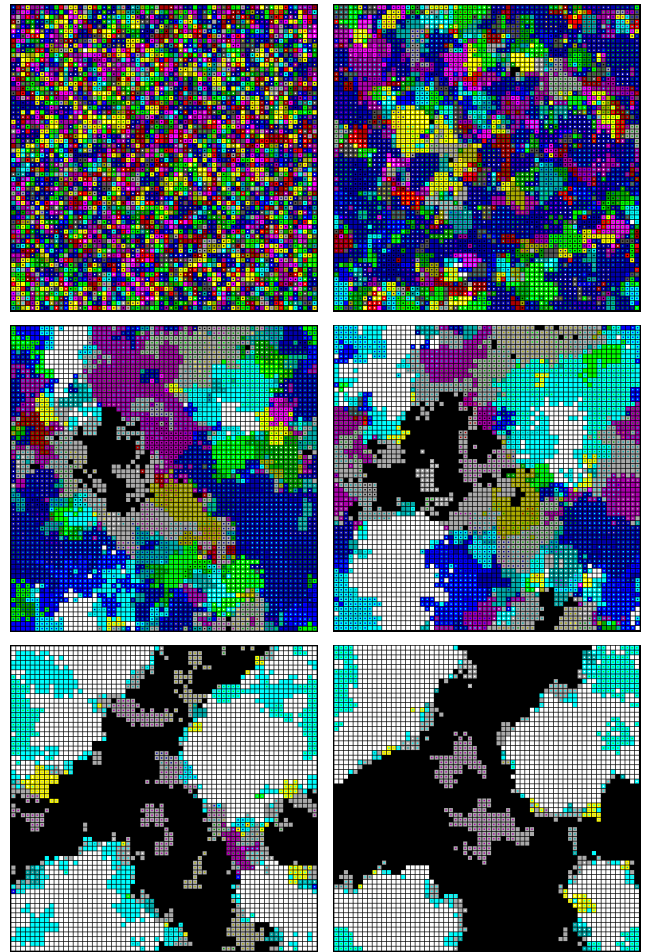


Figure 3 Emergence of two dialects of Communicators, shown in solid black and white, in a randomized array of neural nets with partial training on successful neighbors.

Comparison Agents: Imitators, Localized Genetic Algorithms, and Neural Nets

Here, for the sake of simplicity, our environments contain wandering food sources but no 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 can remain silent. Mouth opening carries an energy cost of .95 points, with an energy cost of .05 points for sounding.

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, of which four are of particular note. 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. Strategy $\langle 0,0,1,0 \rangle$ is a ‘free rider’; it opens its mouth when it hears a sound, benefiting from a signal from a Communicator neighbor, but does not signal reciprocally when fed. The null strategy $\langle 0,0,0,0 \rangle$ does nothing—it never opens its mouth and never makes a sound—and so pays no energy costs. ‘All Eat’ $\langle 0,0,1,1 \rangle$ keeps its mouth open constantly, harvesting any food that comes by but never making a sound.

It should be noted that we use ‘imperfect’ worlds throughout. All cells follow their programmed strategies subject to a 5% measure of error. Nowak and Sigmund (1990) argue 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 and Kilb, 2000).

Although our sample space of behaviors is the same across our studies, those behaviors are instantiated in different ways—as coded behaviors or as operating neural nets. This allows us to compare strategy change by imitation, by localized genetic algorithm, and by local training of neural nets side by side.

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 neighbor’s strategy 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, should a cell have a more successful neighbor, its 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. It should be noted that ours is a *localized* genetic algorithm. All genetic recombination is local: cells with locally successful neighbors change their strategies to local hybrid recombinations.

In a third series of runs we generate the same sample space of behaviors using very simple neural nets (Figure 4). For simplicity, we use bipolar inputs with weights and biases ‘chunked’ at one-unit intervals between -3.5 and $+3.5$. If the sum at the output node exceeds a threshold of 0, the output is treated as $+1$, and the individual opens its mouth, for example. If less than or equal to 0, the output is treated as -1 , and the individual keeps its mouth closed.




Figure 4 Simple neural nets

For our neural nets, strategy change is by partial training on successful neighbors. With bipolar coding and within the limits of our value scale, using ‘target’ for the neighbor’s output, we can calculate the delta rule as simply $w_{new} = w_{old} + (\text{target} \times \text{input})$ and $\text{bias}_{new} = \text{bias}_{old} + \text{target}$.

Though their behavior ranges are identical, our agents instantiate three very different forms of updating mechanisms. What we want to compare is their behavior across a range of different environments.

Environmental Variability and the Emergence of Communication


1. Communication in a Constant Environment

Our constant environment contains exactly 50 food sources each generation, but we use the gain allotted for successful feeding as an independent variable: tests are run with gains from 1 to 140 points for each successful feeding. We plot what strategy an array evolves to—Communicators or otherwise—and in what number of generations.


Figure 5 shows results across our three modes of strategy change. Runs are to 1500 generations, with the height of each bar indicating how many generations were required to fixation on a single strategy. Should no single strategy occupy the entire array by 1500 generations, the bar tops out, showing the strategy dominant in the array at that point.

It is immediately obvious, and somewhat surprising, how large the window for communication is in each of these cases. Communicators dominate the array from the

Imitation in a Constant Environment




Localized Genetic Algorithm in a Constant Environment




cases All Eat continues to prove dominant above a gain of 90 or 100.

Imitation in a Random Environment



Neural Nets in a Constant Environment



Localized Genetic Algorithm in a Random Environment




Figure 5 Conquest by All Eat at gains of 90 and above using strategy change by imitation, by localized genetic algorithm, and by partial neural net training on successful neighbors.

case in which each successful feeding is worth 10 points to the case in 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 proves dominant.

2. Communication in a Random 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 6 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 gain in the upper limit for Communicators. In all

Neural Nets in a Random Environment




Figure 6 Window for communication in an environment of randomized food sources with a mean of 50. Conquest by All Eat in each case at the upper end.

3. Strong Emergence of Communication in a 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 in the variability of food resources? 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 patterns of ‘boom and bust’ may vary greatly. 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 2 (Sin+2), we get one pattern of numbers for our food sources. With values taken at intervals of 3 (Sin+3), we get a different series (Figure 7).

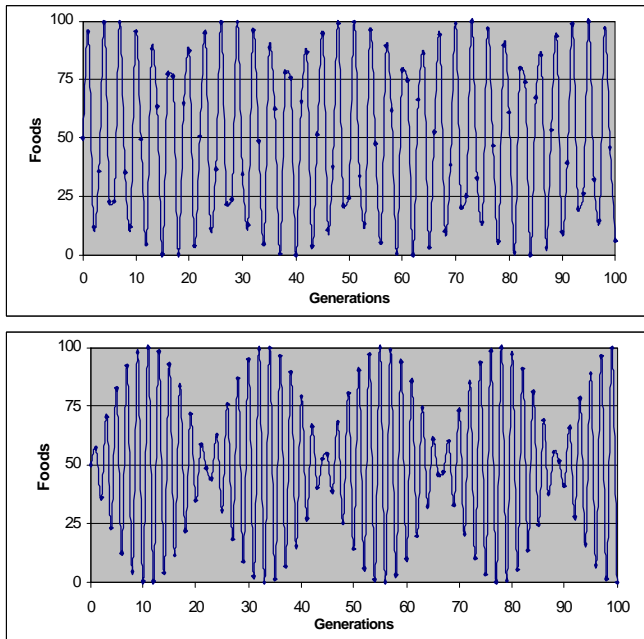


Figure 7 Different patterns of variation in food items: $[\sin(x) + 1] * 50$ for different incremental series x_0, x_1, \dots, x_n . In the Sin+2 series at top, $x_{n+1} = x_n + 2$. In the sin+3 series below, $x_{n+1} = x_n + 3$.


What impact does a sine-wave variable environment have on the emergence of communication? Figure 10 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 our 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 tried larger gains for successful feeding up to 500, beyond the scope of the graph; it is still the Communicators that succeed.


We have found the result to be sensitive to patterns of variability, of course, but it is by no means confined to the pattern of sin+2. Resources following the pattern of sin+3 show similar results.

Across all of our modes of strategy change, then, sine-wave variable environments show a dramatic widening of the window of gain values in which Communicators appear and flourish. Although the average number of food

Imitation in a Sin+2 Variable Environment



Localized Genetic Algorithm in a Sin+2 Environment



Neural Nets in a Sin+2 Environment




Figure 10. Triumph of Communicators at all gains above 10 in a ‘boom and bust’ environment on the pattern of sin+2, for all forms of strategy change.

sources remains the same as in our constant and randomly variable environments, an environment of 'boom and bust' strongly favors the emergence of communication. Although we have focused on simpler studies here, more complicated environments involving both food sources and predators show a similar effect.

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, 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, and Kokalis 2002).

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 which they are constant?

For an environment with randomly variable resources, the answer is 'no'. Random variation shows 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.

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 (Chesson and Huntly, 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 and higher primates, with speculative links to social learning (Potts 1996, Opdyke 1995, Boyd & Richerson 2000). The impact of environmental variability on individual learning is perhaps most developed in decades of careful work on schedules of reinforcement (Nye 1992). We take it as a suggestive fact, worthy of further investigation, that a clear impact of one form of environmental variation is evident even in simulations as simple as those we have outlined here.

References

- Boyd, R., and Richerson, P. 2000. Built for Speed: Pleistocene Climate Variation and the Origin of Human Culture. In Tonneau, F. and Thompson N. eds. *Perspectives in Ethology*, vol. 13: Evolution, Culture, and Behavior. New York: Kluwer Academic/Plenum, pp. 1-45.
- Chesson, P. L. and N. Huntly, 1997. The Roles of Harsh and Fluctuating Conditions in the Dynamics of Ecological Communities. *American Naturalist* 150: 519-553.
- 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., St. Denis, P., & Kokalis, T. 2002. Learning to Communicate: The Emergence of Signaling in Spatialized Arrays of Neural Nets. *Adaptive Behavior* 10: 45-70.
- Honig, W. and Staddon, J. eds. 1977. *Handbook of Operant Behavior*. Englewood Cliffs, NJ: Prentice Hall
- Nowak, M., and Sigmund, K. 1990. The Evolution of Stochastic Strategies in the Prisoner's Dilemma. *Acta Applicandae Mathematicae* 20: 247-265.
- Nye, R. 1992. *The Legacy of B. F. Skinner: Concepts and Perspectives, Controversies and Misunderstandings*. Pacific Grove, CA: Brooks/Cole.
- Opdyke, N. 1995. Mammalian Migration and Climate over the Past Seven Million Years. In Vrba, E., Denton, G., Partridge, C., and Burckle, L. eds. *Paleoclimate and evolution with emphasis on human origins*. New Haven, CT: Yale Univ. Press. Pp. 8-23.
- Potts, R. 1996. *Humanity's Descent: The Consequences of Ecological Instability*. New York, William Morrow.
- Reynolds, R. G. 1986. An Adaptive Computer Model for the Evolution of Plant Collecting and Early Agriculture in the Eastern Valley of Oaxaca. In Flannery, K. ed. *Guilá Naquitz: Archaic Foraging and Early Agriculture in Oaxaca, Mexico*. New York: Academic Press, pp. 439-500.