The Emergence of Communication through Synthetic Evolution*

Technical Report UT-CS-99-431

Bruce J. MacLennan

Computer Science Department
University of Tennessee, Knoxville
maclennan@cs.utk.edu

October 20, 1999

Abstract

We describe four series of experiments to study the emergence of *inherently meaningful communication* by synthetic evolution in a population of artificial agents, which are controlled either by finite state machines or by neural networks. We found that the agents can evolve the ability to use single symbols and, to a limited extent, pairs of symbols exhibiting rudimentary syntax. We show that the communication system evolved by the population can be studied in its own right as an evolving emergent phenomenon, and that the emergent communication systems exhibit some of the richness of natural communication.

1 Introduction

This article describes a series of experiments to study the emergence of *inherently meaningful communication* by synthetic evolution in a population of artificial agents. By "inherently meaningful" we mean that the communication is meaningful and relevant to the agents themselves, independent and regardless of any meanings we (as observers) may attribute to the communications. (We discuss elsewhere [7, 13] the relevance to the study of intrinsic intensionality of these experimental techniques, which

*To appear in *Advances in Evolutionary Synthesis of Neural Systems*, edited by Vasant Honavar, Mukesh Patel, and Karthik Balakrishnan (MIT Press).

we call *synthetic ethology.*) Briefly, we may say that communication is inherently meaningful if it has some actual or potential relevance to the agents.

However, Burghardt [1] has defined communication as "the phenomenon of one organism producing a signal that, when responded to by another organism, confers some advantage (or the statistical probability of it) to the signaler or his group." Therefore communication acquires its primary, natural meaning through the selective advantage it has conferred on the communicators through their evolution. (Hence, ecologically valid investigations of communication must take place in the agents' natural environment or something close to it.) Thus we may conclude that meaningful communication must be investigated in an evolutionary context, and that if inherently meaningful communication is to emerge in artificial systems, it will do so in a context of *synthetic evolution*.

2 Experimental Design

2.1 Synthetic World

If we are to design an synthetic world in which genuine communication can be expected to emerge, then we must begin by asking what sort of world will make this likely. For communication to emerge, it must have some selective advantage, such as allowing the agents to coordinate their behavior more effectively.

One way to accomplish this is to give a reproductive advantage to agents that coordinate their behavior, for example, when the behavior of one agent A is appropriate to a *situation* known only to the second agent B. We can think of this *situation* either as being some aspect of B's internal state (so it is not observable by A), or as being some external situation perceivable by B but out of A's range of perception. Of course, these analogies are ultimately irrelevant; all that matters are the formal structures of perceivability and selection. Since our goal has been to design an experiment that is as simple as possible while still exhibiting the phenomenon of interest (communication), these conditions of coordination have been established in a simple way.

These conditions are created by giving each agent a *local environment*, the state of which is perceivable by that agent but by no other. Of course the state can be simple or complex, but in these experiments we have kept it very simple. Furthermore, although one could use any sort of process to determine the states of the local environments, we set them randomly, which makes it as difficult as possible for an agent's local-environment state to be predicted by the other agents.

Our goal is to select for coordinated behavior among the agents and our hope is that communication will emerge as a way of accomplishing this. Therefore we select for an action that corresponds to the local environment of another agent. That is, if the local environment of agent B is in state λ, and a different agent A acts in a way appropriate to λ, then we say that A and B have cooperated and we give them some reproductive advantage. In accord with our goal of keeping the experiment as simple as possible, we have interpreted "acting in a way appropriate to λ" to be a simple match with λ. That is, A has succeeded in cooperating with B when it manages (by whatever means) in matching the state of B's local environment.

There are many ways in which such cooperation could be made to affect reproductive advantage; we have taken the simple approach of awarding a point of abstract credit to each agent that cooperates every time it cooperates. Agents are more likely to reproduce if they have accrued more credit and they are more likely to die if they have accrued less. Cooperation need not be limited to pairs of agents, and so in general if several agents $A_1, \ldots, A_n$ match B's local environment state, then all $n+1$ agents will be awarded credit (n points for the emitter B, 1 point for each of the actors $A_1, \ldots, A_n$. Other variations include imposing credit penalties for unsuccessful attempts to cooperate (i.e. mistaken attempts to match the other's local environment state).

Clearly, communication will never emerge if it is physically impossible in the simulated world. In particular there must be some potential medium of communication and the agents must have the physical capability to alter and sense the state of that medium. (We say "must" because in these experiments the mechanisms of synthetic evolution are incapable of evolving new sensory or effector organs, so we must take these as givens. We do not, however, give them any specific function, such as communication.)

Therefore, in these experiments we provide the synthetic world with a *global environment* whose state can be sensed and altered by all the agents. By "sensed" we mean that an agent's behavior may depend on the state of the global environment, and by "altered" we mean that an agent's behavior can modify the state of the global environment. As with the local environments, the state of the global environment may be simple or complex in structure, but in the experiments described below it has been kept as simple as possible. Because communication takes place in time, for example by a signal that varies through time, we must allow temporal variations of the global environment's state; that is, the agents can cause it to change through time and can sense those variations through time.

As a consequence of the foregoing, we can see that the agents must have a behavioral control mechanism with the following faculties:

1. ability to respond to the state of its local environment,
2. ability to respond to the state of the global environment,
3. ability to alter the state of the global environment,
4. internal memory or other mechanism for responding to and influencing the time-course of the global state,
5. ability to attempt cooperative activity with other agents.

There are many mechanisms that can satisfy these requirements, including recurrent neural networks, finite state machines, traditional programs, and rule-based systems such as classifier systems.

One of our goals has been to study the evolution of language, and one of the characteristics of language is that it is a cultural phenomenon transmitted through learning. Therefore, in some of our experiments we allow our agents to learn, which means that there is some means for automatic adaptation of their behavioral control mechanisms. Reinforcement, unsupervised and supervised learning are all possible mechanisms for behavioral adaptation.

2.2 Synthetic Evolution

As remarked in section 1, ecological validity requires that inherently meaningful communication be investigated in its "natural" environment, that is, in the environment in which it has evolved. This implies that the emergence of communication must occur in the context of synthetic evolution, which allows the behavioral mechanisms of the agents to evolve in accord with the selective pressures of the synthetic world. Our basic approach is straight-forward.

Each agent has a genetic string that encodes its genotype. When an agent is created this string is used to create the agent's phenotype, which represents its behavioral mechanism. In some cases the agent's phenotype can change by means of learning or other forms of adaptation.

Recall that agents are awarded credit for successful cooperation; this credit is used to influence the probability of an agent reproducing or dying. In our experiments reproduction is sexual. Two agents are chosen as parents, with reproductive preference given to agents that have accrued more credit. Their genetic strings are combined by a simplified model of biological crossover with a low probability of random mutation; the result becomes the genetic string of the offspring.

Agents may also die, which means that their phenotypes are removed from the population; agents with less credit are more likely to die than those with more credit. In the experiments to be described we have kept the population size constant with a one-for-one replacement rule: an agent must die for each that is born.

2.3 Data Collection

We collect various kinds of data to track the emergence of communication and to study its evolving structure. Since we select for cooperation and cooperation is (by design) difficult in the absence of communication, an important variable is the amount of cooperation in the population. Therefore, we measure the average number of cooperations in a given interval of time (i.e., the average credit accrued in that interval); this quantifies the *degree of coordination* in the agents' behavior.

We are especially interested in the emergence of communication as a collective behavior of the population, that is, we want to study the emergent communication as manifested by the population as a whole and evolving through time. To understand this we can, for example, measure correlations between the apparent signals and external (local environment) states and between them and the internal (memory) states of the agents. In this sense we (as observers) can discover the meanings created by the agents for their (inherently meaningful) signals.

There are several statistics that may be computed. For example, if the global environment is being used sometimes as a medium for communicating the local-environment state, then there should be a correlation between local- and global-environment states when successful cooperations take place. In the absence of communication there should be no systematic relation between global- and local-environment states. One way to quantify this is to count the fraction of times each pair occurs simultaneously with a cooperation. The resulting probability matrix has maximum entropy (reflecting its lack of order), when the signals have no systematic meaning, and the entropy decreases as the signal use becomes more regular. This is one way in which we can quantify the emergence of a communication system. Another statistic of deviation from randomness is the coefficient of variation, which is the standard deviation in units of the mean ($V = \sigma/\mu$).

We may also gather relevant statistics of the population of agents that supports the communication system. For example, to better understand the structure of the populations we may compute statistics on the weight matrices of neural nets or on the transition tables of finite state machines.

2.4 Experimental Controls

One of the advantages of studying communication through synthetic evolution is the degree of experimental control that it permits. I will mention a few examples.

Certainly, when compared with the study of communication in nature, one of the advantages is that we can have complete control of the genetic structure of the population. For example, we can run two simulations under different conditions with genetically identical populations, or we can seed the population with selected genotypes for particular experimental purposes. In addition, since learning can be enabled or disabled, we can compare situations in which the agents are able to learn or not, or in which they have different adaptive mechanisms or parameters. Further, if some interesting phenomenon appears in a simulation, we can rerun it and make interventions for the sake of the investigation.

One form of control that we have found especially useful is the external suppression of the possibility of communication: by frequent randomization of the state of the global environment (effectively raising its noise level) we can prevent it from being used for communication and thereby prevent communication from emerging. This permits comparison of population behavior under conditions in which communication can and cannot evolve.

Finally, one of the advantages of our approach is that the behavioral mechanism is completely available for analysis. At any point during the simulation we may subject any individual agent or the entire population to analysis, thus relating the internal behavioral mechanisms to the structure and evolution of the emergent communication system.

3 Series 1: One-symbol Communication by FSMs

3.1 Setup

The first series of experiments investigated the simplest form of single-symbol communication. The local-environment state λ was drawn from a small discrete set Λ of size $L = |\Lambda|$. In most of these experiments $\Lambda = \{0, \ldots, L-1\}$, where $L = 8$. Likewise, the global environment state γ was drawn from a discrete set Γ of the same size, $G = |\Gamma| = L$. In practice, $\Gamma = \{0, \ldots, 7\}$.

The behavioral control mechanism was a finite state machine (FSM) with one internal state (and hence no memory). We chose finite state machines because they are a simple behavioral model that has the potential of both generating and recognizing sequential signals (although that potential was not exploited in the first series of experiments). Thus, each machine's transition table had GL entries for each possible combination of global- and local-environment states.

The finite state machine can respond in only two ways: to *emit* (or *signal*) by altering the global-environment state or to *act* by attempting to match another machine's local-environment state. (Recall that a machine's local environment is not under its control.) In effect, each table entry represents one of two kinds of rules. An *emission rule* has the form:

$$(\gamma, \lambda) \Longrightarrow \mathrm{emit}(\gamma'),$$

where (γ, λ) is the current global/local state and $\mathrm{emit}(\gamma')$ makes γ' the new (altered) global state. Similarly, an *action rule* has the form:

$$(\gamma, \lambda) \Longrightarrow \mathrm{act}(\lambda'),$$

where $\mathrm{act}(\lambda')$ attempts to match λ' to the local environment of another machine. Thus a machine has $G + L$ possible responses, encoded as an integer in the range $0, \ldots, G + L - 1$.

Observe that a machine's response always depends on both the global state and its local state. This means that its response to a given signal is context-dependent, for it is potentially different in each local situation in which it may find itself. Therefore, it will not automatically respond to a signal in the same way in all situations, although it may evolve to do so; therefore the machines face a difficult evolutionary challenge (more on this in section 6.1).

The genotype of a machine is simply represented by a string of GL genes, each with $G + L$ alleles. In these experiments there were 64 genes with 16 alleles. Two-point crossover was used, which means that two numbers η, θ were chosen randomly from $1, \ldots, GL$. The genetic strings were treated like rings; that is, between η and θ the offspring's genes were copied from one parent, and between θ and η from the other. With low probability (0.01) a single randomly selected gene was mutated to a random allele. Population size was kept constant by having the offspring replace a low-scoring agent.

It is necessary to mention a modification to the rules for cooperation: we judge a cooperation to have taken place only if an agent's action matches the local-environment state of *the last emitter*. The reason is that with only eight possible local-environment states and modest population sizes (e.g. 100), it would be almost certain that any action would match the local environment of *some* other agent. Therefore, cooperation would be easy by "guessing" and there would be little selective pressure toward the emergence of communication. Even with this more restrictive cooperation rule there is a 1/8 chance of guessing correctly without communication. (Further consequences of this cooperation restriction are discussed under "Partial Cooperation" in section 3.3.)

The process of synthetic evolution is organized into three nested cycles. The outermost are the *breeding cycles*, in each of which one agent is chosen to die and two agents are chosen to breed (producing a single offspring). Each of the B breeding cycles comprises $E = 10$ *environmental cycles*, at the beginning of each of which, all the local environments are set randomly. (Thus the states of the local environments cannot be predicted.) Each environmental cycle comprises $A = 5$ *action cycles*, during each of which all the agents have an opportunity to respond to the global- and local-environment states. In these experiments the agents are serviced cyclically (a decision discussed in section 7).

The probability of being selected as a parent was proportional to accrued credit (number of cooperations) in a breeding cycle, while the probability of dying was inversely related in a simple way to accrued credit (the exact formulas given elsewhere [6, 7]).

We investigated a very simple single-case learning rule for these FSMs. When a machine attempts unsuccessfully to cooperate, its transition table is altered to give what would have been the correct response in these circumstances. That is, if under conditions (γ, λ) the machine responded $\mathrm{act}(\lambda')$, but the local environment state of the last emitter was $\lambda'' \neq \lambda'$, then the (γ, λ) entry of the transition table is altered to be $\mathrm{act}(\lambda'')$.

We ran our simulations under three different conditions: (1) communication suppressed, (2) communication permitted with learning disabled, and (3) communication permitted with learning enabled. By "communication suppressed" we mean that a random signal was written into the global environment after each agent responded, thus preventing the possibility of communication through the global environment. By

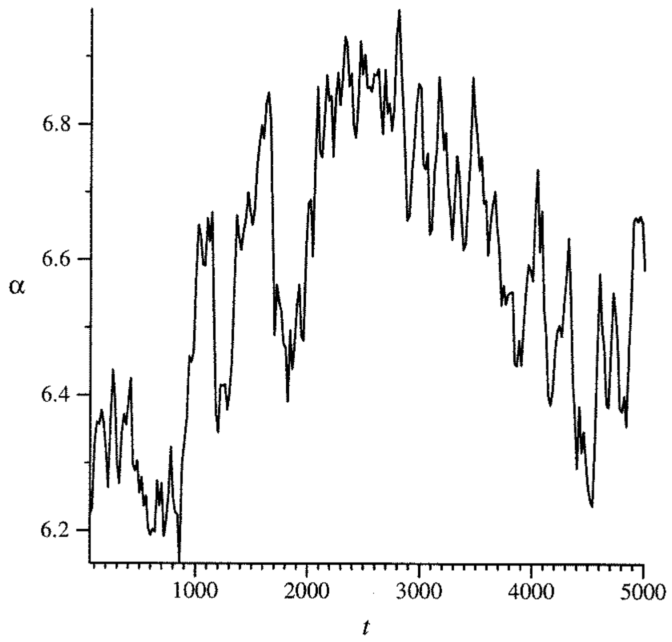

Figure 1: Degree of Coordination α: Communication Suppressed

"communication permitted" we mean that the global environment was not randomized in this way; however, we do nothing directly to facilitate or encourage its use for communication. By this control we can measure the selective advantage of the emerging communication system.

By "learning enabled" or "disabled" we mean that the previously described learning rule is or is not allowed to operate. In the former case the phenotype can diverge from that determined by the genotype, in the latter it cannot. There are of course many other sorts of controls that can be used with these experiments, but even these few generate interesting phenomena.

In these experiments, the population size was $P = 100$. The simulations were usually run for $B = 5000$ breeding cycles, although some simulations were run for much longer.

3.2 Results

We ran a series of more than 100 simulations of this kind; in most cases we ran genetically identical random starting populations under all three conditions. As is to be expected from simulations of this kind, there is considerable variation from run to run, but all the results are qualitatively the same as those we will describe. The experiments are robust and have been replicated in other laboratories [16].

When communication is suppressed, the degree of coordination (average level of

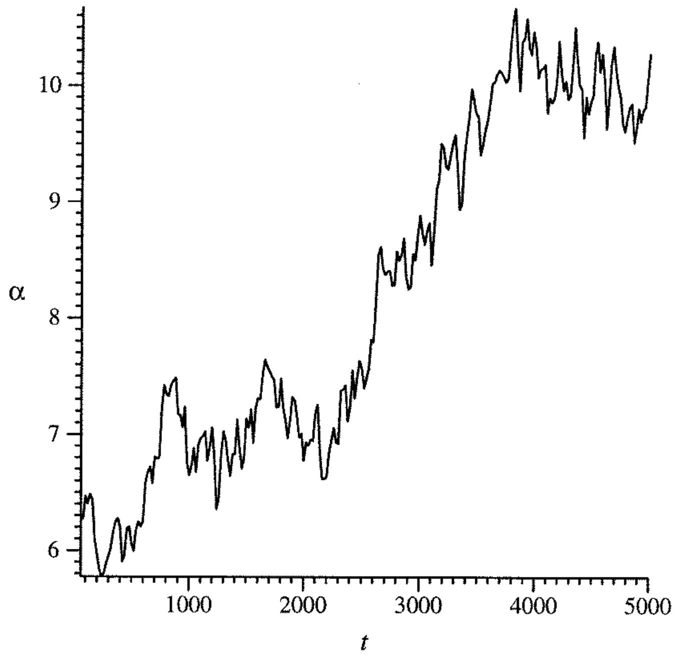

Figure 2: Degree of Coordination α: Communication Permitted with Learning Disabled

cooperation) stays near to 6.25, the calculated level when the agents are "guessing" [6]; figure 1 shows how the average number of cooperations per breeding cycle α varies over time (measured in breeding cycles). Although we would expect the degree of coordination to stay near the chance level, a linear regression analysis shows a slight upward trend, 3.67×10^{-5} cooperations/breeding cycle/breeding cycle. This is a stable phenomenon, which will be discussed in section 3.3.

Figure 2 shows the evolution of the degree of coordination α, as measured by average cooperations per breeding cycle, when communication is not suppressed. It is apparent to the eye that coordination is increasing much faster than when communication was suppressed, a fact confirmed by linear regression, for the degree of coordination is increasing at a rate of 9.72×10^{-4} cooperations/breeding cycle/breeding cycle, which is 26 times faster than it was when communication was suppressed. After 5000 breeding cycles the average number of cooperations per cycle has grown to 10.28, which is 65% above the level of 6.25 achievable without communication.

In figure 3 we see the result when communication is not suppressed and the agents are able to learn from their mistakes. First, it is apparent that communication starts at a much higher level than under the two previous condition. This is because after making a mistake an agent has four more opportunities in an environmental cycle to respond correctly before the local environments are re-randomized. Further, the degree of coordination is increasing much more rapidly than without learning: $3.71 \times$

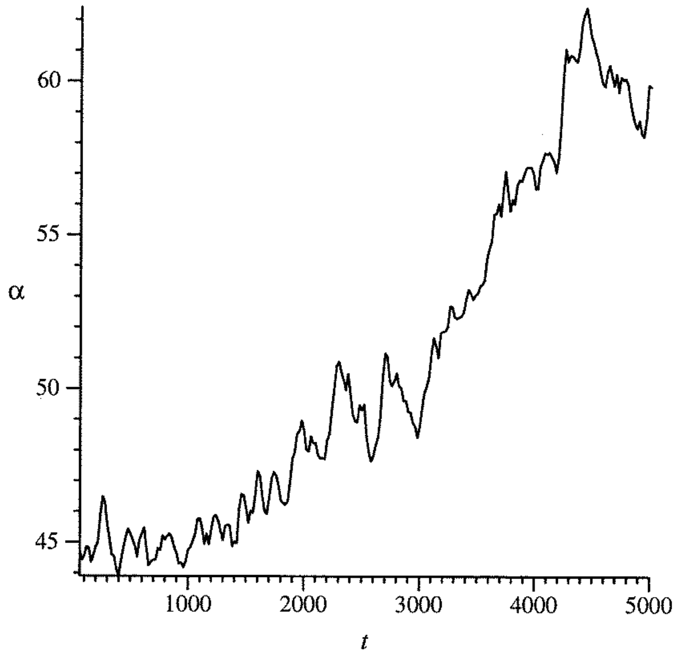

Figure 3: Degree of Coordination α: Communication Permitted with Learning Enabled

10^{-3} cooperations/breeding cycle/breeding cycle, which 3.82 times the rate without learning and 100 times the rate when communication was suppressed. After 5000 breeding cycles, the degree of coordination has reached 59.84 cooperations/cycle, which is 857% above the level achievable without communication.

The preceding results show us *that* communication has emerged and that it has significant selective advantage, but it does not tell us much about the structure of the emerging communication system. As suggested in section 2.3, we can keep track of the co-occurrence of local- and global-environment states that occur together in successful cooperations. That is, whenever a cooperation takes place we increment a count corresponding to the state of the global environment and the state of the local environment of the last emitter. We cannot be sure from such a co-occurrence that the global-environment state *means* the local-environment state, but nonrandom associations between the two will point in that direction.

Since in these experiments we are most interested in the later stages of the evolution, we calculate the co-occurrence tables over the last 50 breeding cycles of the simulation.

Table 1 shows the co-occurrence table that resulted when communication was suppressed.[1] Although some structure is apparent, overall the local and global states

[1] Crumpton [2, App. A] discovered a small error in the calculation of the co-occurrence matrix in

Table 1: Co-occurrence Matrix: Communication Suppressed

	situation							
sym.	0	1	2	3	4	5	6	7
0	94	130	133	34	166	0	150	682
1	16	105	279	228	261	307	0	118
2	0	199	229	12	0	0	161	274
3	95	19	93	283	669	89	0	201
4	1	97	212	200	112	0	0	0
5	28	135	84	8	600	215	0	351
6	0	0	0	118	59	70	0	690
7	0	33	41	0	371	0	0	0

Table 2: Co-occurrence Matrix: Communication Permitted with Learning Disabled

	situation							
sym.	0	1	2	3	4	5	6	7
0	0	0	2825	0	500	20	0	0
1	206	0	0	505	999	231	2	0
2	1	0	0	277	39	4935	1	2394
3	385	1	1	94	0	0	1483	1
4	0	292	0	0	19	555	0	0
5	0	0	1291	0	0	144	0	0
6	494	279	0	403	0	1133	2222	0
7	140	2659	0	202	962	0	0	0

are weakly correlated. The entropy is $H = 4.95$ bits, which is lower (more ordered) than the calculated maximum entropy $H_{\max} = 6$ bits (derivations given elsewhere [6]). The coefficient of variation is $V = 1.27$. These numbers will be more meaningful when we compare them to the other two conditions.

Now consider the co-occurrence matrix that results when communication is not suppressed (table 2). The table is visibly more organized than when communication was suppressed. This is confirmed by the coefficient of variation $V = 2.13$, which is larger than in the suppressed case, $V = 1.27$, reflecting a less-random use of signals. For comparison, $V_{\min} = 0$ for a uniform co-occurence matrix and $V_{\text{ideal}} = \sqrt{7} \approx 2.65$ for an "ideal matrix" [7], which has a one-to-one correspondence between local- and global-environment states. The entropy $H = 3.87$ bits, which is closer to the entropy $H_{\text{ideal}} = 3$ bits of an ideal co-occurence matrix than it was when communication was

the communication-suppressed case, which made it appear less structured than it is. Table 1 reflects the corrected calculation; the corresponding tables in prior publications [6, 7, 14] are incorrect. Noble and Cliff [16, table 2] also noted the discrepancy.

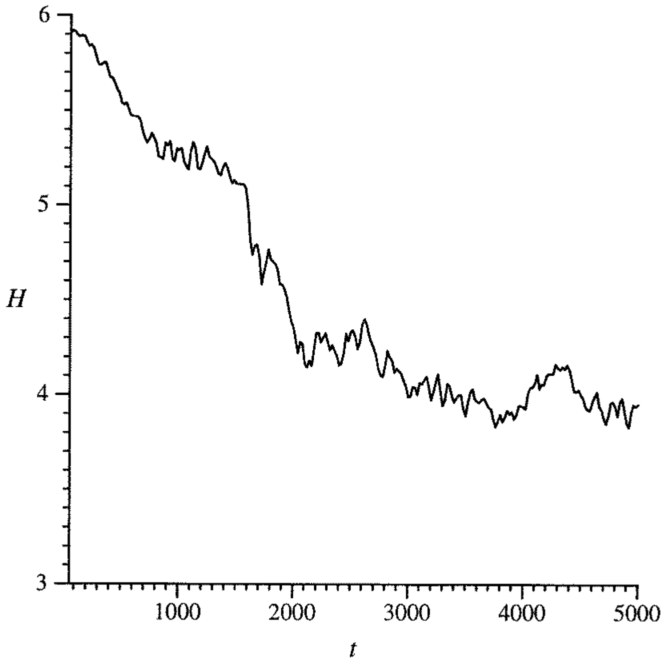

Figure 4: Entropy: Communication Permitted with Learning Disabled

suppressed ($H = 4.95$ bits). Figure 4, which shows the change of entropy over time, demonstrates the emergence of an ordered communication system in the population.

Table 3 shows the co-occurrence matrix that resulted when communication was not suppressed and the agents were able to learn. The coefficient of variation is $V = 2.39$, which is a little larger than in the non-learning case; the entropy is $H = 3.91$ bits, which is a little larger than in the non-learning case. This is fairly typical: the entropy with learning may be a little larger or smaller than without it.

Table 2 shows some of the richness typical of natural communication. For example symbol $\gamma = 7$ is most often associated with situation $\lambda = 1$ and vice versa, although it sometimes denotes situation 4, and situation 1 is occasionally represented by symbols 4 and 6. There are also cases of synonymy, for example situation 6 may be represented by symbols 3 or 6. Further, we find ambiguity, for example, symbol 4 may represent situations 1 or 5 (or occasionally 4). Such synonymy and ambiguity could result from individual agents using synonymous or ambiguous symbols, from the existence of competing dialects in the population, or from a combination of the two, but experiments by Noble and Cliff (discussed in section 7) point to the first possibility.

Finally there is asymmetry in symbol use. For example, situation 7 is virtually always denoted by symbol 2, which is however ambiguous, and more commonly denotes situation 5. Similarly, symbol 5 almost always denotes situation 2, which is

Table 3: Co-occurrence Matrix: Communication Permitted with Learning Enabled

	situation							
sym.	0	1	2	3	4	5	6	7
0	3908	29172	1287	12281	2719	1132	93	3836
1	191	634	107	1039	0	0	2078	0
2	4675	1306	0	37960	85	410	7306	26611
3	0	410	0	0	0	126	1306	304
4	0	0	353	62	575	1268	420	519
5	36	0	46	469	0	0	0	26
6	1075	156	0	0	0	951	0	1086
7	0	73	54	0	2764	135	461	102

Table 4: Summary of Order Measures

		Communication/Learning			
Measurement	Random	N/N	Y/N	Y/Y	Ideal
Coefficient of Variation, V	0	1.27	2.13	2.39	2.65
Entropy, H (bits)	6	4.95	3.87	3.91	3

however mostly denoted by symbol 0.

The values of the entropy and coefficient of variation in each of the three conditions, along with their extreme values, are collected in table 4. Overall it is apparent that not suppressing communication allows the emergence of an organized communication system, regardless of whether the agents are capable of learning.

3.3 Partial Cooperation

We must pause to consider a phenomenon we call *partial cooperation* (inaccurately termed "pseudo-cooperation" in some earlier reports) [7, 14]. Recall (section 3.1, p. 7) that we have placed a restriction on cooperation — an actor must match the local state of the last emitter — since otherwise chance cooperations will be much too easy. However, this restriction creates a loophole in the scoring algorithm, which the population may evolve to exploit. Specifically, the agents may coadapt to emit and act in only a subset ($\Lambda' \subset \Lambda$) of the local-environment states. This strategy raises the chances of a correct guess to $1/|\Lambda'|$ from $1/|\Lambda|$. The evolution of the population to exploit the loophole explains the slow increase in the degree of coordination when communication is suppressed (figure 1); it also explains why in long simulations the agents communicate about a decreasing subset of the situations. It is *genuine cooperation*, but occurs without the benefit of communication by restricting the opportunities for potential cooperation. In spite of these difficulties, we kept the cooperation restriction, since it facilitated the emergence of communication in shorter simulations. (See section 7 for Noble and Cliff's investigation of this phenomenon.)

4 Series 2: Two-symbol Communication by FSMs

4.1 Setup

A second series of experiments was intended to investigate the possibility of finite-state machines evolving the ability to generate and recognize sequential signals (sequences of global-environment states used for communication). We accomplished this by creating an artificial world in which the number of local-environment states is greater than the number of global-environment states, so a single symbol (global-environment state) cannot uniquely specify a situation (local-environment state). That is, since $G < L$ there is no map from Γ onto Λ, although there are maps from Γ^* onto Λ.

We decided to begin with local environments that could be expressed by two symbols; that is, there are maps from Γ^2 onto Λ, so $G^2 \geq L$. In this case we chose $G = 4$ and $L = 8$, so two symbols are more than enough to express the local states.

Obviously, if the agents are to be able to recognize or generate sequential signals, they must have some memory by means of which to control their sequential behavior. Therefore, in this series of experiments the agents were finite state machines (FSMs) with $S = 4$ possible internal memory states. That is, we gave them the minimum memory necessary to remember a global-environment state ($S = G$). Let Σ be the internal state space, so $S = |\Sigma|$; in practice $\Sigma = \{0, \ldots, S-1\}$.

These machines are defined in effect by a set of behavioral rules of the form

$$(\sigma, \gamma, \lambda) \Longrightarrow (\sigma', R)$$

for all possible $\sigma \in \Sigma, \gamma \in \Gamma, \lambda \in \Lambda$. In practice, the rules are represented by a transition table of SGL entries, indexed by $(\sigma, \gamma, \lambda)$. Each table entry contains the pair (σ', R), where $\sigma' \in \Sigma$ is the new internal state and R is a response, either emit(γ') or act(λ'), as before. Thus, a table entry must encode $S(G+L)$ possibilities. In practice we represent this as a pair of numbers, one in the range $0, \ldots, S-1$, the other in the range $0, \ldots, G+L-1$. Similarly, the genotype is represented by a string of SGL genes, each chosen from $S(G+L)$ alleles. In these experiments there were 128 genes with 48 alleles.

Except for the difference in genetic structure, the mechanism of synthetic evolution was essentially the same as in the first series of experiments (section 3). However, we did try several variations of the selection strategy, such as imposing a penalty for failed attempts to cooperate and making the probability of selection proportional to the square of the number of cooperations per breeding cycle. The nested breeding, environmental and action cycles were also similar, except that the simulations were generally run longer: 10^4 to 2×10^5 breeding cycles. Learning was implemented by

Table 5: Co-occurrence Matrix: Communication Permitted with Learning Disabled

	situation							
sym.	0	1	2	3	4	5	6	7
0/0	31	22	42	0	144	0	0	0
1/0	26	15	62	0	175	0	0	0
2/0	119	23	44	0	47	0	0	0
3/0	8	9	18	0	31	0	0	0
0/1	0	54	106	2	74	59	516	0
1/1	0	33	174	3	423	227	1979	0
2/1	0	23	65	17	139	74	125	0
3/1	0	1	24	0	48	96	51	0
0/2	50	4	4	366	7	0	8	42
1/2	35	9	0	32	1	0	6	44
2/2	52	76	0	112	7	0	13	135
3/2	52	6	1	215	2	0	2	78
0/3	0	2	13	17	0	3	0	0
1/3	0	66	19	6	0	4	0	0
2/3	0	33	61	27	0	2	0	0
3/3	0	39	38	8	0	0	0	0

a similar single-case algorithm: if the rule $(\sigma, \gamma, \lambda) \Longrightarrow [\sigma', \mathrm{act}(\lambda')]$ was applied but the correct action was λ'', then the rule is changed to $(\sigma, \gamma, \lambda) \Longrightarrow [\sigma', \mathrm{act}(\lambda'')]$. For experimental purposes, as before, communication can be suppressed by randomizing the global environment state, and learning can be enabled or disabled. Further, the population size was $P = 100$ and the mutation rate was 0.01, as in the first series.

4.2 Results

Table 5 shows the co-occurrence matrix resulting from a typical simulation, which ran for 10^4 breeding cycles (communication unsuppressed, learning disabled). When a successful cooperation takes place, we increment the table entry corresponding to the local-environment state (column) and to the *last two* global-environment states (row). There is obvious structure in the matrix. For example, the table falls into 4×4 submatrices of similar degrees of coordination, which means that successful cooperations tend to be more sensitive to the most recent of the two symbols, rather than to the first of the two. For example, local-environment state 5 is usually expressed by signals of the form $X1$ (that is 01, 11, 21, or 31). This suggests that the machines are not making full use of their memory capacity. Nevertheless, the agents sometimes make full use of the expressive power of two symbols. For example, 00, 10 and 30 usually mean local state 4, but 20 usually means local state 0, so here the machines

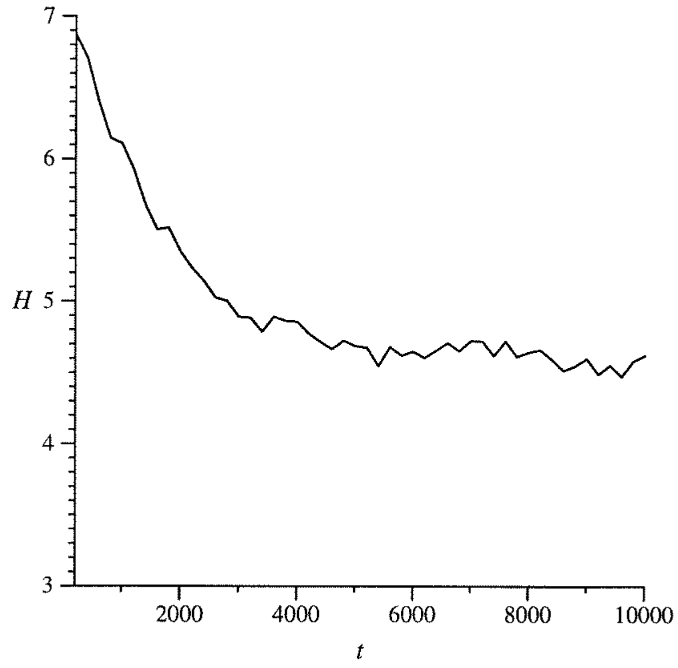

Figure 5: Entropy of Two-symbol Communication (Learning Disabled)

are using the first symbol to modify the meaning of the second. Furthermore, order is significant, since 02 usually denotes local state 3 and only occasionally 0.

Nevertheless, in this series of experiments we never observed the agents evolving to make full use of their communicative potential. We can see this in figure 5, which shows the decrease of entropy H as communication evolves; it begins at $H_{\text{max}} = 7$, reflecting total disorder, and decreases to about $H = 4.5$, which is still substantially more disordered than the $H_{\text{ideal}} = 3$ of a "perfect" two-symbol communication system. The entropy appears to have stopped decreasing after about 5000 breeding cycles, so we can see that longer simulations are not likely to help (nor did they, in our experiments).

It appears that two-symbol communication cannot fully self-organize, at least under the conditions investigated in these experiments. We can understand why by considering the task that the agents must evolve the ability to solve. Recall that in each action cycle, all of the machines are allowed to respond as determined by the global environment γ, their local environment λ and their internal state σ. The machines are serviced cyclically, which means that once a machine emits a symbol, it will have to wait while all the other machines are serviced before it has an opportunity to emit a second symbol. Consider what must take place for a machine A to signal its local environment λ^* to another machine B by means of the symbol sequence $\gamma_1\gamma_2$. First, whatever the internal state σ of machine A and whatever the state γ of the global environment, machine A must have a rule $(\sigma, \gamma, \lambda^*) \Longrightarrow [\sigma_1, \mathrm{emit}(\gamma_1)]$. Second, supposing that no machine between A and B has altered the global environment, whatever the internal state σ' of machine B and whatever the state λ' of its local environment, machine B must have a rule $(\sigma', \gamma_1, \lambda') \Longrightarrow [\sigma'', R]$. If R is $\mathrm{act}(\lambda')$, then (a) it is attempting to cooperate prematurely, and will succeed only if $\lambda' = \lambda^*$ by chance. On the other hand, if R is $\mathrm{emit}(\gamma')$, then (b) it will change (unless perchance $\gamma' = \gamma_1$) the global environment, destroying the chances of any other machines responding to A (or of B doing so on later action cycles). Third, machine A must have a rule $(\sigma_1, \gamma'', \lambda^*) \Longrightarrow [\sigma_2, \mathrm{emit}(\gamma_2)]$, where γ'' is the global state resulting from the last machine to emit before the second servicing of A (as conditioned by its own internal and local states and by the global state). Fourth, again supposing that no machine between A and B has altered the global state, machine B must have a rule $(\sigma'', \gamma_2, \lambda') \Longrightarrow [\sigma''', \mathrm{act}(\lambda^*)]$. The emergence of this degree of coordination may be too much to expect, and it is perhaps surprising that we observed as much use of two-symbol signals as we did.

5 Series 3: Paired Symbol Emission by FSMs

5.1 Setup

The difficulty of evolving two-symbol communication led my student Joseph Crumpton to consider a modification of the experimental design [2]. Because of the likelihood of two-symbol signals being disrupted by intervening emissions, he decided to give each agent *two* chances to respond in each action cycle. Notice that Crumpton's machines do *not* simply emit a pair of symbols $\mathrm{emit}(\gamma_1\gamma_2)$; rather it is still necessary that the machines use their internal state to control the generation or recognition of sequential signals. In an action cycle each machine is allowed to cycle twice before proceeding to the next machine. The global-environment state comprises two symbols, which we may call the first and second components of the global state. On the first of a machine's cycles it is senses the first component and can change the first component by emission; on the second cycle it senses the second component and can change the second component by emission. On both cycles, the machine's response is, as before, dependent on its internal and local-environment states. Furthermore, in addition to emitting or attempting to cooperate, Crumpton's machines are allowed to do nothing, so that they are not forced to emit or attempt to cooperate while waiting for a second symbol. As in the second series of our experiments (section 4), Crumpton used $S = 4, G = 4, L = 8, P = 100$, but ran the simulation for 6×10^4 breeding cycles. The selection, breeding and learning processes are also the same as in the earlier experiments.

5.2 Results

In a typical simulation run, Crumpton found that the number of cooperations per breeding cycle increased to 65 in the first 10^4 breeding cycles; it remained between 65 and 70 for the remaining 5×10^4 breeding cycles of the simulation, although there was a slight elevation in the average level after $t = 3 \times 10^4$ cycles. This is significantly higher than the 23 cooperations/breeding cycle obtained in the earlier experiments (analysis of variance $p < 0.01$). Similarly, the entropy decreased to about 5.1 bits in the first 10^4 breeding cycles; after a slight elevation between $t = 2 \times 10^4$ and $t = 3 \times 10^4$, the entropy stabilized at $H = 5$ bits for the remaining 3×10^4 cycles. This was *not* significantly different from the entropy achieved in the second series of experiments (section 4).

For each local-environment state, Crumpton considered the symbol pair most commonly used in successful cooperations; such a pair might be said to express that state (perhaps ambiguously), and so we'll refer to it as an *expressive pair*. For example, in a typical run he found five expressive pairs (11, 22, 32, 33, 34), which means that the eight environment states could not be expressed unambiguously. The average he measured (4.3) was not significantly different from that found in the earlier experiments (4.0).

Crumpton was especially interested in non-repeating symbol use, that is, signals of the form XY, which require the use of memory in their generation and recognition, as opposed to signals of the form XX, which do not. For example, of the five expressive pairs (11, 22, 32, 33, 34), two are non-repeating (32, 34). In his experiments Crumpton found an average of 1.4 non-repeating pairs, which is significantly higher ($p < 0.01$) than the 0.3 of the earlier series of experiments (section 4), thus showing that the machines were making better use of the representational resources of the medium.

Crumpton used his experimental design to investigate several other properties of communication in his synthetic world. For example, he found significant differences in entropy and degree of coordination for population sizes $P = 50, 100, 200, 400$ and 800. (The $P = 800$ case was exceptional in many respects, which Crumpton attributes to the population being too big for communication to have stabilized in the 6×10^4 breeding cycles of his experiments.) In general, larger populations achieved a higher degree of coordination, but $P = 100$ led to the most use of expressive and non-repeating pairs. Degree of coordination was found to be significantly negatively correlated with entropy, as would be expected.

Crumpton also investigated differing amounts of memory for the agents ($S = 2, 3, 4, 5, 6$). With regard to degree of coordination and entropy, smaller numbers of internal states were better, but this was found to be deceptive, since with $S = 2$ or 3 the eight situations were being represented by an average of three expressive pairs, whereas with $S = 4$ they were represented by an average of four expressive pairs. That is, in the former cases the agents were achieving higher degrees of coordination by cooperating in a smaller subset of the situations (as was discussed in section 3.3, p. 13).

Crumpton ran a series of experiments in which there were two "species" of agents: the usual ones, and a breed of memoryless "competitors" who received credit for blocking cooperation by matching the last emitter's local environment.[2] It was hoped that the competitors would push the communicators to use pairs of symbols. The result, however, was a significantly lower ($p < 0.01$) degree of coordination (48) and significantly less ($p < 0.03$) use of non-repeating symbols (0.4). He observed that the competitors either dominated the population or had very little effect.

Crumpton investigated a number of other variations (including variations in learning rule and placement of offspring in the population), that did not lead to significant differences in degree of coordination, entropy, or the use of expressive and non-repeating pairs of symbols.

6 Series 4: One-symbol Communication by Neural Networks

6.1 Motivation

There are at least two limitations to the FSM behavioral control mechanisms used in the preceding experiments. First, one of the important problems in the evolution of language is the emergence of discrete signal types from a continuum of states [7, 8, 9, 10, 11]. (Steels has addressed this problem directly as part of his investigation of symbol grounding [17, 18].)

Second, since FSMs have in effect a separate rule (i.e. transition table entry) for each possible combination of internal, global and local state, they have no inherent capacity to generalize. For example (and ignoring internal state for now), to consistently signal γ for situation λ, the machine must have in effect the rule $(*, \lambda) \Longrightarrow \text{emit}(\gamma)$, where "$*$" means "don't care" (i.e. any possible global state). However, conventional FSMs do not permit rules of this form (although alternative rule-based models, such as classifier systems, do). Therefore, the FSM must acquire (through evolution or learning) a rule $(\gamma', \lambda) \Longrightarrow \text{emit}(\gamma)$ for every possible global-environment state γ'. The situation is even worse when the machines have internal state too. This problem is aggravated by the fact that the size of the table, and hence of the genetic strings, increases with the product of the sizes of the state spaces, so there is no economy for the machines in discovering general rules.

We originally made the decision to use FSMs because they include no *a priori* model of general rules, but this same characteristic means that the population has a much more difficult evolutionary problem to solve. Hence, it seemed appropriate to investigate behavioral control mechanisms more capable of representing and learning

[2] That is, if a competitor was able to match the last emitter's local environment, it prevented the signaller and any communicating responders from getting credit, and it received credit for every such blocked communication.

general rules.

6.2 Setup

My students Rick Stroud and Noel Jerke [15, sec. 2] conducted an exploratory investigation patterned after our first series (section 3), but using neural networks instead of FSMs, and continuous rather than discrete local and global state spaces, specifically, $\Gamma = [0, 1] = \Lambda$. Each neural net had two inputs for sensing the states of its local environment and the global environment, and two (continuous-valued) outputs representing the emission/action choice (in $[-1, 1]$) and the emission/action value (in $[0, 1]$). In addition the nets had a single hidden layer comprising six neurons.

The emission/action choice was indicated by the sign of the first output. In the emission case the second output becomes the new state of the global environment. In the action case the second output is compared to the local environment of the last emitter; credit is awarded if they differ by less than 1/8 (in emulation of the $L = 8$ discrete states of the earlier experiments).

The overall cycle of synthetic evolution was the same as in the first series of experiments (section 3). The genetic string represented the connections between the neurons ($+1, -1, 0$; i.e. excitatory, inhibitory, absent) but not their relative strengths, which were adjusted by back-propagation. Generally five to ten cycles of back-propagation were used to train the networks. The emit/act output was trained to ± 1, depending on its sign on its first response (i.e., its native response was strengthened). Further, when the response was an action, the other output was trained to match the last emitter's local-environment state. Simulations were run for 10^4 breeding cycles. The continuous local and global state spaces were each divided into ten bins for compiling a 10×10 co-occurrence matrix.

6.3 Results

Stroud and Jerke observed that communication did not emerge when equal credit was awarded to the emitter and actor for successful cooperations (as was done in the FSM experiments), but that it did emerge when the actor was awarded three units and the emitter one. Without this adjustment, they speculated, there was an increase in the tendency of the agents to emit, and therefore an increased probability of emissions interfering with each other; the better coordinated populations seemed to have a small number of emitters and a large number of actors. With increasing coordination, the corresponding co-occurrence matrices displayed an obvious increase of structure, but Stroud and Jerke did not quantify it by entropy or other measures. Although Stroud and Jerke did observe an increase in the cooperation of the agents, it seems to have resulted in part from the "partial cooperation" of the agents exploiting the loophole in the scoring rule (section 3.3), rather than from full-fledged cooperation, for the co-occurrence matrices show cooperations to be taking place in a decreasing subset of the local state space.

7 Related Work

Noble and Cliff [16] have replicated our first series of experiments (section 3) and extended them in a number of ways. Overall their results agreed with ours, but their experiments exhibited several discrepancies. First, they measured lower entropy in the communication-suppressed condition, as did Crumpton [2] when the program was corrected (footnote 1, page 10). Second, they observed a lower degree of coordination than we did when learning was enabled; this has not been explained. Next, they compared the sequential (cyclic) servicing of the agents that we used with servicing them in a different random order each time, and they found that the results were unaffected except when learning was enabled, in which case coordination increased more quickly with random updating (and in fact agreed with the values we measured). Further, they investigated whether synonymy and ambiguity in the emergent communication system reflects the existence of multiple dialects in the population or ambiguous symbol usage by individual agents (recall section 3.2, p. 12). The latter interpretation was supported, since over time the population tends toward genetic homogeneity. Third, they investigated extreme situations in which there were a small number of possible states ($L \leq G \leq 4$). They found that the agents were able to achieve comparatively high coordination by exploiting the loophole in the scoring rule (section 3.3).

8 Conclusions

It will be worthwhile to summarize our results: (1) We have demonstrated consistently that inherently meaningful communication can emerge through synthetic evolution in populations of simple artificial agents. (2) This has been demonstrated for agents controlled by both finite state machines and neural networks. (3) The agents can evolve the ability to use single symbols and, to a limited extent, pairs of symbols exhibiting rudimentary syntax. (4) The communication system evolved by the population can be studied in its own right as an evolving emergent phenomenon. (5) The emergent communication systems exhibit some of the richness of natural communication, including ambiguous, synonymous, asymmetric and context-dependent symbol use.

We will mention a few opportunities for future work. In nature, communication media are inherently continuous, but some functions of communication are facilitated by the use of discrete symbols. This implies that the emergence of discrete symbols from continuous phenomena is an important problem that must be solved by the evolution of language [7, 8, 9, 10, 11, 14]. As mentioned above, Steels has initiated a promising investigation of this problem [17, 18]. In addition to explaining the emergence of words, such research may illuminate the gradual emergence of the rules and hierarchical structure characteristic of human language [8, 11]

We have already remarked on the limitations of the FSM behavioral model (section 6.1), but the simple neural nets employed by Stroud and Jerke are not much better. Animals, on the other hand, independent of any communicative abilities they may have, are able to interpret perceptually complex phenomena and to generate complex, flexible, hierarchically structured motor behaviors. It seems likely that the underlying neural mechanisms of these behaviors may be recruited and adapted for communication, including language. Therefore, we may speculate that our artificial agents may need to be similarly endowed with perceptual-motor abilities before they will evolve complex, structured communication systems. (For this reason, our recent research has been directed toward mechanisms for perception and motor control [12].)

The experiments we have described involve local environments that are extremely simple in structure: a small discrete set (e.g. $\{0, \ldots, 7\}$) or a simple continuum (e.g. $[0, 1]$); the communication systems that have evolved to describe them are correspondingly simple. Although some of the structure (syntax) of communication systems may be emergent phenomena independent of their semantics and pragmatics [3, 4, 5], we expect the complexity of communication to reflect the complexity of the agents that use it and the complexity of their world. That is, structure in the pragmatic interactions of agents induces structure in the semantics that must be expressed, which induces structure in the syntax to express it. This suggests that we will have to construct more structured synthetic worlds in order to study the emergence of more structured communications systems, including language.

References

[1] Gordon M. Burghardt. Defining 'communication'. In J. W. Johnston Jr., D. G. Moulton, and A. Turk, editors, *Communication by Chemical Signals*, pages 5–18. Appleton-Century-Crofts, New York, 1970.

[2] Joseph J. Crumpton. Evolution of two symbol signals by simulated organisms. Master's thesis, University of Tennessee, Knoxville, December 1994.

[3] Simon Kirby. Language evolution without natural selection: From vocabulary to syntax in a population of learners. Technical Report EOPL-98-1, University of Edinburgh, Department of Linguistics, 1998. Available from `http://www.ling.ed.ac.uk/~simon`.

[4] Simon Kirby. Learning, bottlenecks and the evolution of recursive syntax. In E. J. Briscoe, editor, *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge University Press, Cambridge, 2000. Also available from `http://www.ling.ed.ac.uk/~simon`.

[5] Simon Kirby. Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners. In C. Knight, M. Studdert-Kennedy, and J. Hurford, editors, *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Forms*. Cambridge University Press, Cambridge, 2000. Also available from `http://www.ling.ed.ac.uk/~simon`.

[6] Bruce J. MacLennan. Evolution of communication in a population of simple machines. Technical Report CS-90-99, University of Tennessee, Knoxville, Department of Computer Science, January 1990. Available from `http://www.cs.utk.edu/~mclennan`.

[7] Bruce J. MacLennan. Synthetic ethology: An approach to the study of communication. In Christopher G. Langton, Charles Taylor, J. Doyne Farmer, and Steen Rasmussen, editors, *Artificial Life II: The Second Workshop on the Synthesis and Simulation of Living Systems*, pages 631–658. Addison-Wesley, Redwood City, CA, 1992. Santa Fe Institute Studies in the Sciences of Complexity, proceedings Vol. X. Report available from `http://www.cs.utk.edu/~mclennan`.

[8] Bruce J. MacLennan. Characteristics of connectionist knowledge representation. *Information Sciences*, 70:119–143, 1993. Also available from `http://www.cs.utk.edu/~mclennan`.

[9] Bruce J. MacLennan. Continuous symbol systems: The logic of connectionism. In Daniel S. Levine and Manuel Aparicio IV, editors, *Neural Networks for Knowledge Representation and Inference*, pages 83–120. Lawrence Erlbaum, Hillsdale, NJ, 1994. Also available from `http://www.cs.utk.edu/~mclennan`.

[10] Bruce J. MacLennan. Image and symbol: Continuous computation and the emergence of the discrete. In Vasant Honavar and Leonard Uhr, editors, *Artificial Intelligence and Neural Networks: Steps Toward Principled Integration*, pages 207–24. Academic Press, New York, 1994. Also available from `http://www.cs.utk.edu/~mclennan`.

[11] Bruce J. MacLennan. Continuous formal systems: A unifying model in language and cognition. In *Proceedings of the IEEE Workshop on Architectures for Semiotic Modeling and Situation Analysis in Large Complex Systems*, pages 161–172, Monterey, CA, August 1995. Also available from `http://www.cs.utk.edu/~mclennan`.

[12] Bruce J. MacLennan. Field computation in motor control. In Pietro G. Morasso and Vittorio Sanguineti, editors, *Self-Organization, Computational Maps and Motor Control*, pages 37–73. Elsevier, 1997. Also available from `http://www.cs.utk.edu/~mclennan`.

[13] Bruce J. MacLennan. Synthetic ethology: A new tool for investigating animal cognition. In Colin Allen, Marc Bekoff, and Gordon Burghardt, editors, *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*. MIT Press, forthcoming.

[14] Bruce J. MacLennan and Gordon M. Burghardt. Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2(2):161–187, Fall 1993. Also available from `http://www.cs.utk.edu/~mclennan`.

[15] Bruce J. MacLennan, Noel Jerke, Rick Stroud, and Marc D. VanHeyningen. Neural network models of cognitive processes: 1990 progress report. Technical Report CS-90-125, University of Tennessee, Knoxville, Department of Computer Science, December 1990.

[16] Jason Noble and Dave Cliff. On simulating the evolution of communication. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson, editors, *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, Cambridge, 1996. MIT Press.

[17] Luc Steels. Constructing and sharing perceptual distinctions. In M. van Someran and G. Widmer, editors, *Proceedings of the European Conference on Machine Learning*, Berlin, 1997. Springer-Verlag.

[18] Luc Steels. The synthetic modeling of language origins. *Evolution of Communication*, 1(1):1–34, 1997.