Evolution of Communication in a Population of Simple Machines CS-90-99

Bruce MacLennan

Department of Computer Science
University of Tennessee
Knoxville, TN 37996-1301
maclennan@cs.utk.edu

January 1990

Abstract

We show that communication may evolve in a population of simple machines that are physically capable of sensing and modifying a shared environment, and for which there is selective pressure in favor of cooperative behavior. The emergence of communication was detected by comparing simulations in which communication was permitted with those in which it was suppressed. When communication was not suppressed we found that at the end of the experiment the average fitness of the population was 84% higher and had increased at a rate 30 times faster than when communication was suppressed. Furthermore, when communication was suppressed, the statistical association of symbols with situations was random, as is expected. In contrast, permitting communication led to very structured associations of symbols and situations, as determined by a variety of measures (e.g., coefficient of variation and entropy). Inspection of the structure of individual highly fit machines confirmed the statistical structure. We also investigated a simple kind of learning. This did not help when communication was suppressed, but when communication was permitted the resulting fitness was 845% higher and increased at a rate 80 times as fast as when it was suppressed. We argue that the experiments described here show a new way to investigate the emergence of communication, its function in populations of simple machines, and the structure of the resulting symbol systems.

1 Introduction

1.1 Investigating Communication

What is communication? How can it emerge by natural selection? What form will it take? What are the factors that influence its emergence or form? How do signs come to have meaning? These are some of the questions to which this investigation is addressed.

We believe that these questions will not be answered by armchair theorizing; we seem to have achieved all we can by that approach. Nor will they be answered by studying small populations communicating in unnatural laboratory environments; this approach loses ecological validity since it radically alters the pragmatics of communication. It was the mistake of behaviorism. Therefore, it seems that these questions can only be answered by empirical investigation of populations communicating in their natural environments (or laboratory environments faithful to them in the relevant ways); this of course is the ethological approach.

We do not believe, however, that ethology will answer all our questions about communication. This is because many of the deepest problems pertain to the mental phenomena that accompany the external behavior: When can a symbol be said to have a meaning? What is intentionality? We have argued [9] that answering questions such as these will require an understanding of the neural *mechanisms* by which communication is generated. Thus a complete theory of communication must integrate an ethological account of its function with a neurophysiological account of its mechanism.

Unfortunately, the integrated approach that we envision is beyond the empirical capabilities of contemporary ethology and neuroscience. For this reason, and because science usually progresses fastest when it can investigate phenomena in their simplest and most controllable contexts, we have been studying the evolution of communication in populations of simple machines.

This solves the problems we have mentioned in the following ways. First, the mechanism is transparent. Since these are simple machines (e.g., finite state machines or simple neural networks), we can give a complete account of any communication that takes place. Second, by allowing the population of machines to evolve, we know that if communication emerges then it must confer some selective advantage on the machines that communicate – that is, it is relevant for survival. Thus the pragmatics of communication is preserved, since it is occurring in its "natural" environment (the simple, but complete world of the machines). By this approach we combine ecological validity with the kind of experimental control that has produced the best examples of science. But it is predicated upon our getting populations of simple machines to communicate.

1.2 What is Communication?

How can we tell if two machines are communicating? This is not a trivial question, as shown by the fact that it is sometimes difficult to determine whether or not a given animal interaction is communication. For example, there is more to communication than one organism doing something that another organism notices, since by that definition almost all behavior is communication, and the term looses its significance. We might claim that communication necessarily involves the *intent* of the signaler to influence the receiver's behavior, but if attributing intent to lower animals is controversial, attributing it to simple machines is reckless.[1] Gordon Burghardt's definition of communication seems to provide a way out of this dilemma:

> Communication is 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. [4, p. 16]

In other words, to identify an event as a *communication act* we need to be able to establish: (1) that an organism caused some change in the environment, (2) that a second organism responded to that change, and (3) that the event tends to confer selective advantage on the first organism or its group.

In the case of our simple machines, establishing (1) and (2) is unproblematic (we can simply look at the structure of the machines). Establishing (3), however, requires us to determine the selective value of certain behaviors. The most reliable way of accomplishing this is to follow the evolution of the population, and observe which behaviors confer selective advantage.

This is the approach we have taken in these investigations. We start with populations of randomly generated simple machines (essentially finite state machines) that have the capability of altering the environment in a way that can be sensed by the other machines. Further, we stipulate that the fitness of machines depends on their ability to cooperate, and that each machine's fitness determines its probability of breeding or dying. Thus there is selective pressure in favor of cooperation, but not directly in favor of communication. Our hypothesis is that under these conditions communication – as defined by Burghardt – will emerge. It does.

2 Procedure

2.1 Environment

2.1.1 Structure

There are two components to the environment, a *global environment* and a set of *local environments* (one per machine). All the machines have access to the global

[1] Further, since we hope that these investigations will shed some light on the nature of intentionality, it would be ill-advised to take it as a given.

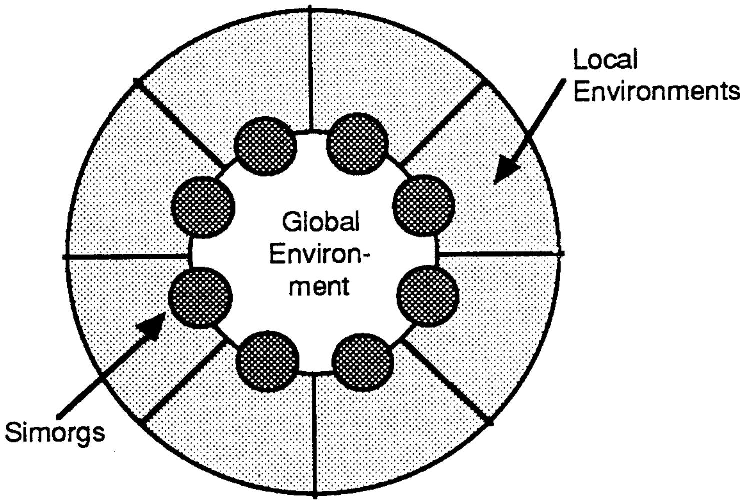

Figure 1: Structure of the Environment

environment, which they can either read or write. On the other hand, each machine can read only its own local environment; it has no direct access to the states of the other local environments. Further, the local environments cannot be written by any of the machines; they are set by an independent random process. (See Fig. 1 for the structure.) Note that the only way one machine can tell the state of its local environment to another is by putting some information about it in the global environment.

2.1.2 States

The contents of the environments are objects drawn from two alphabets. There are G possible global environment states, sometimes called *symbols*, and L possible local environment states, sometimes called *situations*. In this implementation the states are represented by numbers in the range $0, \ldots, G-1$ for the global environment, and $0, \ldots, L-1$ for the local.

2.1.3 Change in the Local Environments

The local environments of all the machines are set at intervals called *minor cycles* or "days."[2] In the current simulation these changes are random, so there is no way they can be predicted by the machines.

[2]We refer to the various time units of the simulation as "hours," "days" and "weeks." These are not intended to have any relation to real hours, days or weeks. They are simply convenient terms for keeping the size of the units in order. Quotation marks are used to remind the reader of this metaphorical use of terminology.

2.2 Machines

2.2.1 Structure and Behavior

The machines that make up the population are a kind of finite-state machine. The state transition is determined by three factors: the machine's current internal state (s), the global environment state (γ), and the machine's local environment state (λ). The result of the transition is a new internal state (s') and a *response*. There are two kinds of responses, *actions* and *emissions*. An emission $\mathsf{emit}(\gamma')$ includes an object γ' drawn from the set of global environment states, and results in this object becoming the new global environment state. An action $\mathsf{act}(\lambda')$ includes an object λ' drawn from the set of local environment states. Such an action is considered "effective" only if it matches the local environment of the last emitter (see Section 2.3). Thus the transition function of each machine is a total function of the form:

$$(s, \gamma, \lambda) \longmapsto (s', r(x))$$

where $r = \mathsf{emit}$ or act and $x = \gamma'$ or λ', respectively.

2.2.2 Learning

The machines can also operate in a mode which permits a simple kind of learning. When learning is permitted, a machine may change its transition table so that in the future it will act in the way that *would have been* appropriate this time. Specifically, suppose that in a context (s, γ, λ) a machine responds $\mathsf{act}(\lambda')$ but that the effective action would have been λ'' (i.e., λ'' was the last emitter's local environment state). Then, learning occurs by replacing the (s, γ, λ) entry of the machine's transition table with $(s', \mathsf{act}(\lambda''))$. Therefore, in the future the transition table will implement

$$(s, \gamma, \lambda) \longmapsto (s', \mathsf{act}(\lambda''))$$

This will hold until the next time that λ'' is an ineffective action in the context (s, γ, λ).

2.2.3 Representation

Since the internal state and local and global environment states are all represented by natural numbers, the transition table can be represented as an array indexed by (s, γ, λ):

$$\begin{array}{c|c|} \cline{2-2} \vdots & \vdots \\ (s, \gamma, \lambda) & (s', r(x)) \\ \vdots & \vdots \\ \cline{2-2} \end{array}$$

Similarly, since s' and x are represented by natural numbers and r can be represented by a bit, the table entries can be represented by triples of natural numbers (see Table 15, p. 31, for an example). The initial population of machines is obtained by setting their transition tables to random numbers in the appropriate ranges.

2.3 Fitness and Breeding

As is common in genetic algorithms and simulations of natural selection, an individual's probability of breeding is *directly* dependent on its "fitness," and its probability of "dying" is *inversely* dependent on its "fitness." This "fitness" is simply a score computed on the basis of the individual meeting some criteria ("acting well"). First we discuss the determination of a machine's fitness, then we discuss the way this influences breeding probability.

2.3.1 Fitness

We want to put selective pressure on the evolution of communication, and one way to do this is to select for cooperative activity that cannot take place reliably without communication. Therefore we consider machines "fit" to the extent that they act in a way that matches *another* machine's local environment. We make the problem harder – and put more pressure in favor of communication – by considering a machine to have acted correctly only when its action matches the local environment *of the last emitter.* In this case both machines, the emitter and the receiver, are credited with a successful cooperation.

Of course, it is quite possible that a machine's action will coincidently match the last emitter's local environment; the frequency with which this can be expected to happen is calculated later (Section 2.7.1). We will be looking for fitness scores above this "guessing" or "chance" level.

Each machine responds a number of times in one "day," at intervals called "hours." The fitness of the machines is accumulated over a longer interval, called a *major cycle* or "week," which comprises a number of "days." Since the local environments change once per "day," the fitness scores reflect each machine's response to a variety of environment states.

2.3.2 Breeding

Once per "week" two individuals are chosen to "breed" and one is chosen to "die" (i.e., be replaced by the single offspring of the breeders). The probability of an organism breeding or dying depends on its fitness score, which has been accumulated over the preceding "week." The probability of breeding is given by

$$p_k = \frac{\phi_k}{P\alpha}$$

where ϕ_k is the fitness of machine k, P is the population size, and α is the average fitness of the population ($\alpha = P^{-1} \sum_j \phi_j$). (If $\alpha = 0$ we set $p_k = 1/P$.) Thus breeding probability is proportional to fitness. We do not require the breeders to be different machines. The probability of dying is given by

$$q_k = \frac{\beta - \phi_k}{P(\beta - \alpha)}$$

where β is the fitness of the most fit individual in the population. (If $\alpha = \beta$ we set $q_k = 1/P$.) Thus probability of dying decreases monotonically with fitness, although not linearly. Since the selection of machines to breed and die is probabilistic, the individual that dies could be one of the breeders.

Each machine has two associated data structures representing transition tables, called the *genotype* and the *phenotype*, which are used in breeding and behavior, respectively. In contrast to genetic algorithms, which usually represent the genotype by a bit string, we represent it by a list containing all the entries in the transition table. For an example, the 64 element list of all the triples shown in Table 15 (p. 31) is a genotype.

The genotypes of the breeders are "crossed over" at two randomly chosen crossover points to yield the offspring's initial genotype (i.e. before mutation). Thus, if $(G_1, \ldots, G_n)$ and $(G'_1, \ldots, G'_n)$ are the parents' genotypes, then the offspring's genotype is

$$(G_1, \ldots, G_j, G'_{j+1}, \ldots, G'_k, G_{k+1}, \ldots, G_n)$$

where j and k are the random crossover points. Crossover is at the level of transition table entries; that is, each gene is a triple $G_i = (s', r(x))$. Thus crossover cannot break up responses; this is different from most genetic algorithms and seems to improve performance in this case.

After crossover, the offspring's genotype is mutated with probability μ. Mutation involves randomly choosing a gene and replacing it with a random triple. Thus we pick random i^*, s^*, r^* and x^* (all in the appropriate ranges), and replace gene G_{i^*} by $(s^*, r^*(x^*))$.

The phenotype is the transition table used to determine the machine's behavior; this is the table in which we look up (s, γ, λ) and which yields the response $(s', r(x))$. In the current experiment the initial phenotype is completely determined by the genotype, since they are both representations of the transition table. Furthermore, if learning is suppressed, they remain the same, since there is no other mechanism for phenotypic change. On the other hand, learning allows the phenotype to change, as discussed above (Section 2.2.2). Notice, however, that it is the genotype that is used for breeding, so there is no genetic mechanism for passing on acquired behavior. (On the other hand, acquired behavior can *indirectly* affect the genotype of the offspring, the so called "Baldwin Effect" [1, 7, 10].)

The noninheritability of acquired behavior leads to an important difference between our breeding algorithm and that common in genetic algorithms. The latter typically implement breeding in distinct "generations," with all of the individuals of the population being replaced at one time. A parent's genetic contribution to the next generation is proportional to its fitness. In contrast, we replace the population incrementally, breeding and killing one machine per "week." Since the probabilities of breeding and dying are determined by fitness, the stochastic behavior of our algorithm should be similar to that typical of genetic algorithms. There is an important exception, however. When learning is permitted, our algorithm permits acquired behavior to be passed from machine to machine, in effect permitting a "culture" to be passed from "elder" machines to "younger" machines (or vice versa!). Since this information is part of the phenotype but not the genotype, the genetic algorithm's wholesale population replacement prevents the information from being passed on (except indirectly through the Baldwin effect). In effect each generation must learn from scratch. We expect such "cultural" transmission to be very important to more sophisticated communicative behaviors. (See also [2, 5].)[3]

2.4 Suppression of Communication

Following Burghardt's definition (p. 3), we will say that communication is taking place only when we can show that some advantage is conferred on the emitter or its group. In this context, this means that communication should be associated with an increased average fitness of the population. But increased relative to what? To determine if communication is taking place it is useful to compare the fitness of the population when communication is possible to that when it is impossible.

To allow this comparison, our simulation has a switch that, when set, *suppresses* communication. This is accomplished by writing a random symbol into the global environment every time a machine responds (regardless of whether the response is an action or an emission). Thus, if any machine is trying to communicate, its attempts will be thwarted, since the global environment is constantly changing outside its control.[4]

When communication has not been suppressed, we say it is *permitted*. This does not mean that communication will take place, only that it will not be actively prevented.

[3]The learning that we implement could still be considered individual rather than social, since a machine's behavior changes only when it acts incorrectly [3]. On the other hand, it would be simple to have machine's learn from other, more fit machines without the cost of their own incorrect action. This fits Boyd & Richerson's definition of culture: "Culture is information capable of affecting individuals' phenotypes which they acquire from other conspecifics by teaching or imitation."[3, p. 33] Imitative learning will be investigated in future experiments.

[4]In an earlier version of this simulation we attempted to suppress communication by replacing the global environment by a random symbol *whenever an emission took place.* The resulting evolution showed that this did not succeed in preventing communication since the machines were still able to communicate (in a limited way) by the presence or absence of symbols!

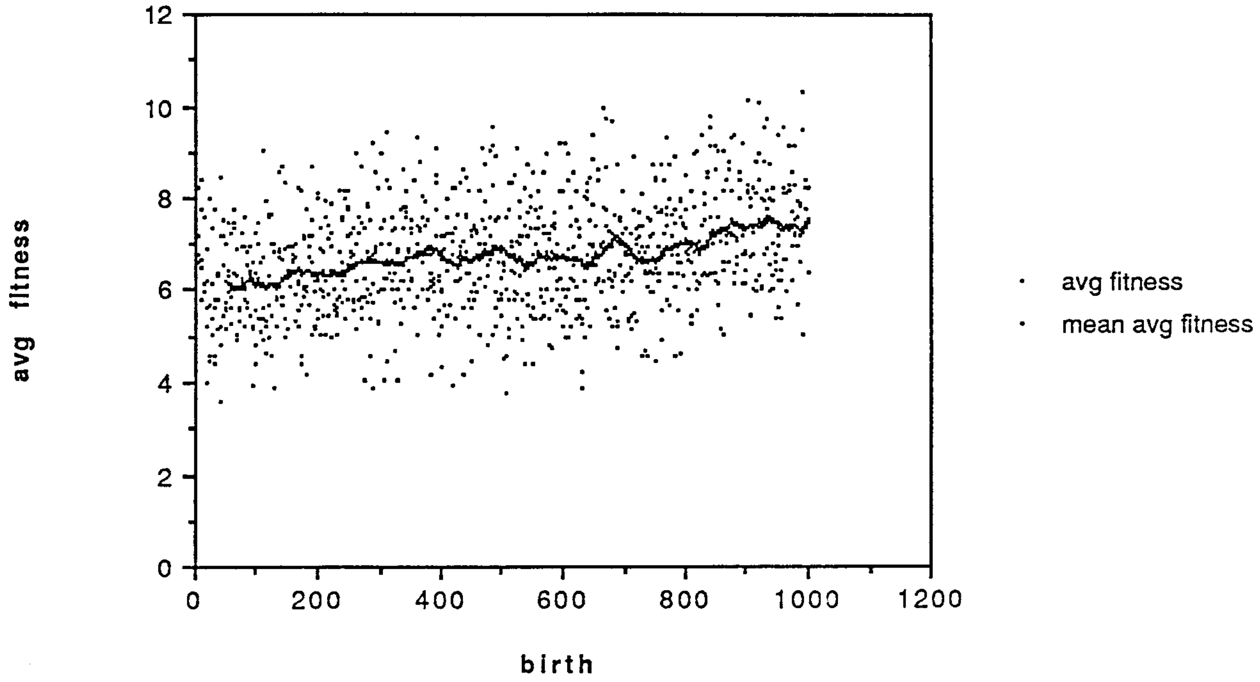

Figure 2: Effects of Smoothing Average Fitness

2.5 Measurements

2.5.1 Fitness

The two fitness parameters we track are the average fitness of the population, and the fitness of the best individual. Since these vary considerably, however, we have found it more useful to analyze a smoothed fitness that results from applying to the raw fitness figures a moving window of size W. Figures 2 and 3 show the effect of smoothing; although there there is still considerable wiggle, a trend is at least visible. In the remainder of this report we refer to smoothed average fitness as α and smoothed best fitness as β.[5] Since the smoothed fitness numbers result from a moving average, the corresponding plot labels are "mean avg fitness" and "mean best fitness."

In comparing evolution under various conditions (e.g., communication suppressed or not, learning permitted or not), it is useful to be able to compare the *rates* at which α and β change, called $\dot{\alpha}$ and $\dot{\beta}$ respectively. To accomplish this we have used linear regression and used the slope of the resulting line as the rate; examples are

[5]A certain quantization is apparent in the best fitness values. The quantum is $H = 10$ "hours"/"day" (in this simulation); since machines act once per "hour" and the local environments change once per "day," fitness often accrues in units this size.

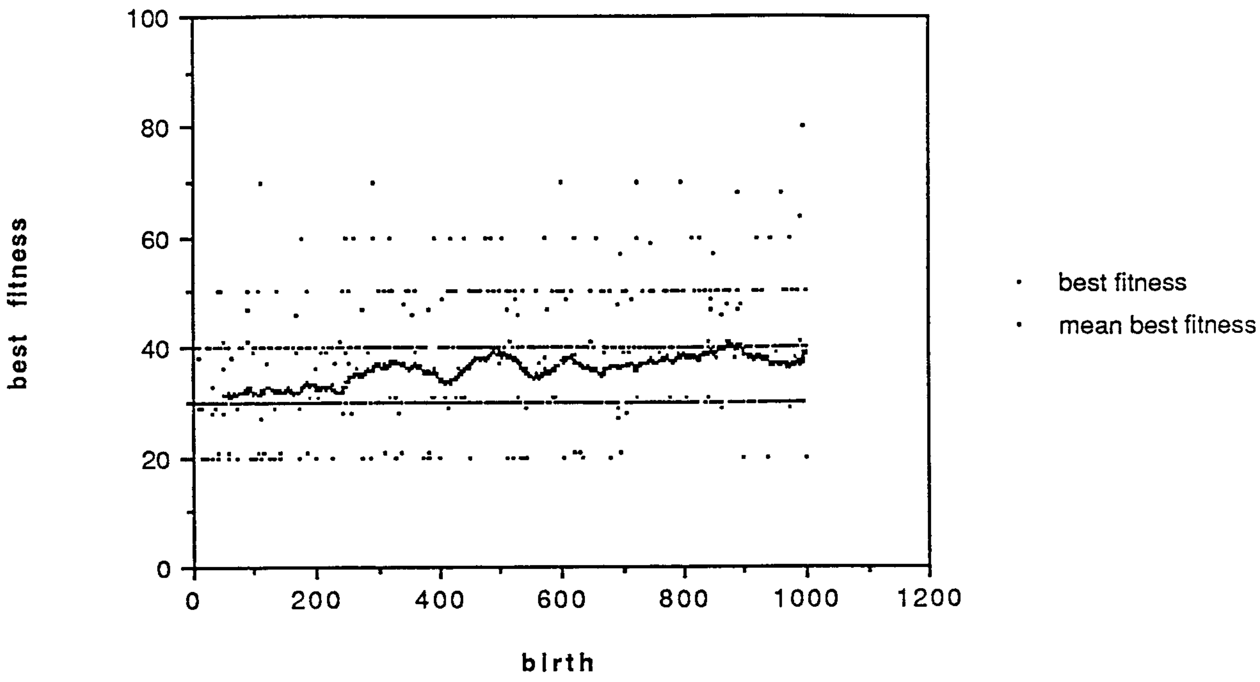

Figure 3: Effects of Smoothing Best Fitness

shown later (Section 3.1). Further, to simplify the computation, we computed the regression only on the plotted α and β values, which were every tenth value in these simulations (see Table 1). Since the regression lines fit rather well, there seems to be no harm in this preprocessing of the data.

2.5.2 Structure of Communication

If the population is evolving the ability to communicate, then this should be apparent in its use of symbols. Therefore, whenever two machines successfully cooperate, we consider an "apparent communication act" to have taken place. (It is only "apparent" because it could have resulted from a lucky guess.) We keep track of such acts and of the symbol/situation pairs they involve. More specifically, suppose that the last emitter put the symbol γ in the global environment and that its local environment is λ. If a later machine responds with $\mathsf{act}(\lambda)$, then they have successfully cooperated, and we say that the second machine *apparently* interpreted symbol γ to denote situation λ. To keep track of this we increment the (γ, λ) entry of a matrix, called the *denotation matrix*.

The only trouble with the foregoing is that early in the simulation symbols will be used randomly, and this random usage may swamp later more organized use. To avoid this, the denotation matrix reflects only those apparent communication acts that occurred in a moving window containing the last W "weeks" of the simulation. Thus the denotation matrices shown in this report reflect only the "recent" use of symbols (see Tables. 3 and 4, pp. 21 and 22, for examples).

If communication is not taking place, and cooperation is being achieved by guessing, then we would expect symbol/situation pairs to occur with equal frequency. The resulting denotation matrix should be very *uniform*, that is, all its entries should be about the same size. On the other hand, we can imagine an "ideal" language to have exactly one symbol for each situation, and vice versa. The resulting denotation matrices would be very *nonuniform*, with only one nonzero entry in each row and in each column. Thus nonuniformity (i.e. deviation from a uniform distribution) reflects structure in the apparent use of symbols to denote situations. How can this structure be quantified?

We have chosen two ways of quantifying the nonuniformity of the denotation matrices. The first makes use of the fact that the standard deviation (σ) of a distribution measures the amount of spread of that distribution around its mean. However, since the actual number of apparent communication acts may differ from run to run, we have to correct for the value of the mean (μ) if we are to get a measure that allows comparisons between runs. Therefore, we use the *coefficient of variation* as a measure of the nonuniformity (structure) of the denotation matrices:

$$V = \sigma / \mu$$

Thus, V measures the amount of spread in units of the mean. For a uniform distribution, $\sigma = 0$ and therefore $V = 0$.

Another measure of the nonuniformity of a distribution is *entropy*. For a two-dimensional discrete probability distribution $p_{\gamma,\lambda}$ the entropy is defined:

$$H = -\sum_{\gamma,\lambda} p_{\gamma,\lambda} \lg p_{\gamma,\lambda}$$

(We use $\lg x = \log_2 x$. We compute the entropy in terms of base 2 logarithms because it gives more meaningful numbers in this case.) The probabilities are computed from the denotation matrix D is the obvious way:

$$p_{\gamma,\lambda} = \frac{D_{\gamma,\lambda}}{\sum_{\gamma,\lambda} D_{\gamma,\lambda}}$$

The maximum state of uniformity has all the probabilities equal, $p_{\gamma,\lambda} = 1/GL$. In this case the entropy is maximum:

$$H_u = -\sum_{\gamma,\lambda} (GL)^{-1} \lg (GL)^{-1} = \lg GL$$

In all the experiments described here, $G = L$, so letting $N = G = L$ we find that the entropy of the uniform distribution is:

$$H_u = 2 \lg N$$

This is the maximum entropy, and represents a completely unstructured use of the symbols.

The minimum entropy occurs when all the entries of D are zero, except one. Such a "δ distribution" has an entropy of 0, which is the minimum.

$$H_\delta = 0$$

This is not a situation we expect to arise, however, since it means that one particular symbol is being consistently used for one particular situation, but that the other symbols are unused and the other situations cannot be denoted. This is an "over-structured" language that's not very useful, so we ask what the entropy would be of the "ideal language," in which one symbol denotes one situation and vice versa. In this case D has N equal, nonzero entries, one in each row and column, which yields an entropy

$$H_I = -\sum_{\gamma}(1/N)\lg 1/N = \lg N$$

Thus the entropy can vary from $2N$, for a completely unstructured language, down to N for the "ideal" language, down to 0 for the "over-structured" language.[6] In these experiments $N = 8$ (Table 1), so $H_u = 6$ and $H_I = 3$.

To simplify interpreting entropies, we introduce a new parameter η representing the lack of structure in the denotation matrix:

$$\eta = \frac{H}{\lg N} - 1$$

We have $\eta = 1$ for the uniform language, $\eta = -1$ for the over-structured language, and $\eta = 0$ for the ideal language.

It is not clear which (if any) of these measures of linguistic structure will prove most useful, so we've listed all three for each denotation matrix shown later.

2.6 Parameters for These Experiments

The parameters used in these experiments are summarized in Table 1. Notice that since the number of internal states is one, the machines have no memory. Hence they are effectively dictionaries that map symbol/situation pairs into responses (either act or emit). They can be visualized as follows:

$$\begin{array}{c|c|}
\vdots & \vdots \\
(\gamma, \lambda) & r(x) \\
\vdots & \vdots \\
\end{array}$$

In this next section we analyze the results that can be expected under these conditions.

[6]Note that an entropy of $\lg N$ does not necessarily mean that the language is "ideal," only that it has the same degree of structure.

Table 1: Parameters Used in Experiments

parameter	value
population	$P = 100$
number of local states ("situations")	$L = 8$
number of global states ("symbols")	$G = 8$
number of internal states	$I = 1$
mutation rate	$\mu = 0.01$
simulation time	$T = 5000$ "weeks"
major cycle	$D = 5$ "days"/"week"
minor cycle	$H = 10$ "hours"/"day"
smoothing window	$W = 50$ "weeks"
plot interval	every tenth "week"
breeding interval	one individual per "week"
fitness interval	accumulated over one "week"
environment changes	once per "day"

2.7 Analysis

2.7.1 Random Population

As a baseline for evaluating the fitness of populations, we estimate the average fitness of a population of random machines (i.e., the contents of the transition tables are completely random). We do this by computing the expectation value of a response, which is

$$\mathcal{E}\{\text{response}\} = \frac{1}{2}\,\mathcal{E}\{\mathsf{act}\} + \frac{1}{2}\,\mathcal{E}\{\mathsf{emit}\}$$

since action and emission are equally likely. If the response is an action, then there is a $1/L$ chance that the action agrees with the local environment of the last emitter. Hence,

$$\mathcal{E}\{\mathsf{act}\} = 1/L$$

If the response is an emission, then the score received by the machine will depend on the responses of the other machines; in particular, it can accumulate points only until the next emission takes place. For each of the following machines there is a $1/2$ chance that it will emit. If it acts rather than emits, then there is a $1/L$ chance that it will act correctly. Thus we can estimate the expectation value of an emission by the infinite series:

$$\begin{aligned}
\mathcal{E}\{\mathsf{emit}\} &= \frac{1}{2}\left(\frac{1}{L} + \frac{1}{2}\left(\frac{1}{L} + \frac{1}{2}\left(\frac{1}{L} + \cdots\right)\right)\right) \\
&= \frac{1}{2L}\left(1 + \frac{1}{2} + \frac{1}{4} + \cdots\right)
\end{aligned}$$

$$= 1/L$$

This is only an estimate, but it is a good one.[7]

Given the foregoing we see that the expectation value for a response is

$$\mathcal{E}\{\text{response}\} = 1/L$$

Therefore, since each machine responds once per "hour," and fitness is accumulated over a "week," which is DH "hours," we see that the expected fitness of a random individual, and hence the average fitness of the population, is:

$$\alpha_r = DH/L$$

For the parameters used in these experiments (Table 1) the fitness of the random population is $\alpha_r = 50/8 = 6.25$.

2.7.2 Optimal Population: Two Varieties

Define a *perfect emitter* to be a machine that in a situation λ always produces a unique symbol γ_λ, and a *perfect receiver* to be a machine that always responds to the symbol γ_λ with $\mathsf{act}(\lambda)$. It is then easy to see that in the optimal population (in terms of average fitness) we have one perfect emitter and $P - 1$ perfect receivers. To understand why, suppose we have a population with E perfect emitters and $P - E$ perfect receivers. When a perfect emitter emits, it will accrue fitness until the next perfect emitter is encountered. If the perfect emitters are distributed uniformly in the population, then we can expect each perfect emitter to accrue $(P - E)/E$ points each "hour" (i.e., each time around the population). In this same period of time, each perfect receiver accrues 1 point. Therefore, since fitness is accumulated over a "week," which is DH hours long, the expected fitness for perfect emitters is

$$\phi_E = DH\frac{P - E}{E}$$

and for perfect receivers is

$$\phi_R = DH$$

Therefore, for a population with E perfect emitters we expect an average fitness

$$\alpha_E^* = \frac{E\phi_E + (P - E)\phi_R}{P} = 2DH\frac{P - E}{P}$$

The expected fitness of the best individual is just ϕ_E:

$$\beta_E^* = DH\frac{P - E}{E}$$

[7] It's an estimate because in fact the series terminates after P terms. The exact expectation value is $(1/L)(1 - 1/2^P)$; in these experiments $P = 100$.

Table 2: Fitness of a Population of Optimal Emitters and Receivers ($P = 100$, $DH = 50$)

E	α_E^*	β_E^*
1	99	4950
2	98	2450
3	97	1617
4	96	1200
5	95	950
10	90	450
12	88	367
20	80	200

Since for there to be both emitters and receivers we must have $0 < E < P$, we see that α_E and β_E are maximized when $E = 1$:

$$\alpha^* = 2DH(P-1)/P, \quad \beta^* = DH(P-1)$$

For later reference, we tabulate in Table 2 α_E^* and β_E^* for several values of E (assuming $P = 100$ and $DH = 50$; see Table 1).

Notice that an optimal population is not very robust; if its single perfect emitter "dies" (which can happen no matter how fit it is), then the average fitness of the population will drop to zero. Thus it seems that only populations with $E \gg 1$ will persist for long; the simulations seem to bear this out (see Section 3.4 and especially Figs. 18 and 19). There is also another evolutionary force towards $E \gg 1$. When $E \approx 1$ the fitness of the perfect emitters is very high relative to that of the receivers (see Table 2). This means that the breeding pair will almost surely be two perfect emitters, and likewise their offspring. Thus the emitters will tend to drive out the receivers (raising E). This will lower the emitters' relative fitness, which will give the receivers a chance to breed. It seems likely that an equilibrium value of E could be reached (although of course the population could have more complicated dynamics).

2.7.3 Optimal Population: One Variety

Another peculiarity of this "perfect" population is that there are two distinct kinds of machines (viz., perfect emitters and perfect receivers). This is also not a stable situation, since a cross of a perfect emitter with a perfect receiver will not produce either. This will become likely as E grows sufficiently to give the emitters and receivers comparable fitness. Thus, it seems unlikely that a perfect population will stay perfect for long. These consideration have led us to impose an additional constraint on the optimal population, namely, that all the genotypes be identical. Since a uniform environment tends to eliminate genetic diversity, this seems a more likely outcome.

To estimate the fitness under this assumption, we first observe that all the machines are identical, and so they can be represented by a single transition table. In each context (γ, λ) this table will produce either an emission or an action. Therefore, suppose that in a given context the machine emits γ. That machine will accrue credit until another machine emits; in the meantime the global environment will remain constant and equal to γ. This means that the actions of the succeeding machines are completely determined by the state of their local environment (which is set randomly) and the γ row of the transition table. Now suppose that in each row there are A actions and $L - A$ emissions. For a perfect population all the actions will be correct, so the score accrued by the emitter will be:

$$\begin{aligned} S &= \frac{A}{L}\left(1 + \frac{A}{L}\left(1 + \frac{A}{L}(1 + \cdots)\right)\right) \\ &= (A/L) + (A/L)^2 + (A/L)^3 + \cdots \\ &= A/(L - A) \end{aligned}$$

This is also the number of receivers involved in these communication acts, and they each receive one point (for a total of S). Since the total number of machines involved is $S + 1$, we compute the average fitness (per hour) as $2S/(S + 1) = 2A/L$. This is maximized when the number of actions per row, A, is $L - 1$ (since otherwise there would be no emissions). Hence, the average fitness of the optimal population is

$$\alpha^* = 2DH(1 - 1/L)$$

For the parameters used in these experiments (Table 1), this is $\alpha^* = 87.5$, which is much closer to the values actually achieved in the simulations (see Fig. 18 and the discussion on p. 32).[8]

3 Results

3.1 Rate of Fitness Change

First we discuss four typical experiments, all of which start from the same population of random machines. Later (Section 3.4) we see how these conclusions generalize to different initial populations.

Figure 4 shows the smoothed average fitness (α) when communication is suppressed and learning is disabled.[9] As expected, α drifts around the chance level

[8] In general, since all the machines are identical, we might expect the best fitness to equal the average. However, by chance a machine might have the opportunity to emit each "hour" in the "week," which would give it a score $\beta^* = DHS = DHA/(L-A) = DH(L-1) = 350$ (for $A = L-1$). This is also in line with the simulations.

[9] Note that in these figures every tenth value is plotted.

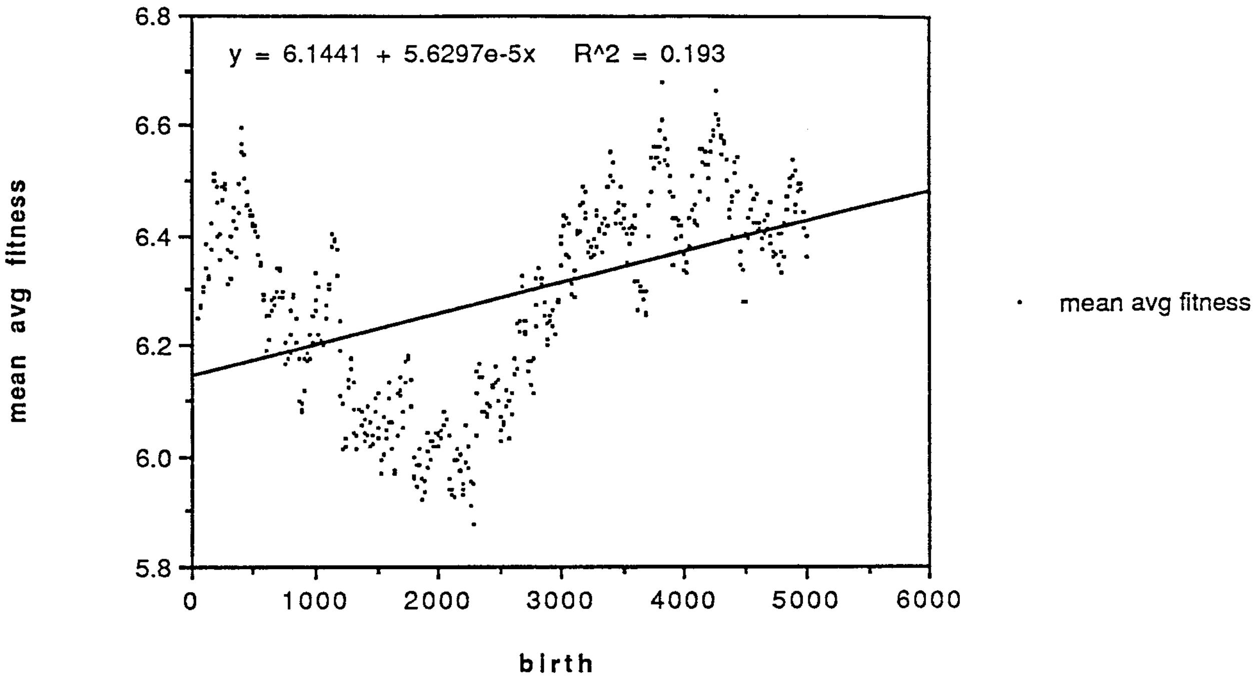

Figure 4: Average Fitness: Communication and Learning Suppressed

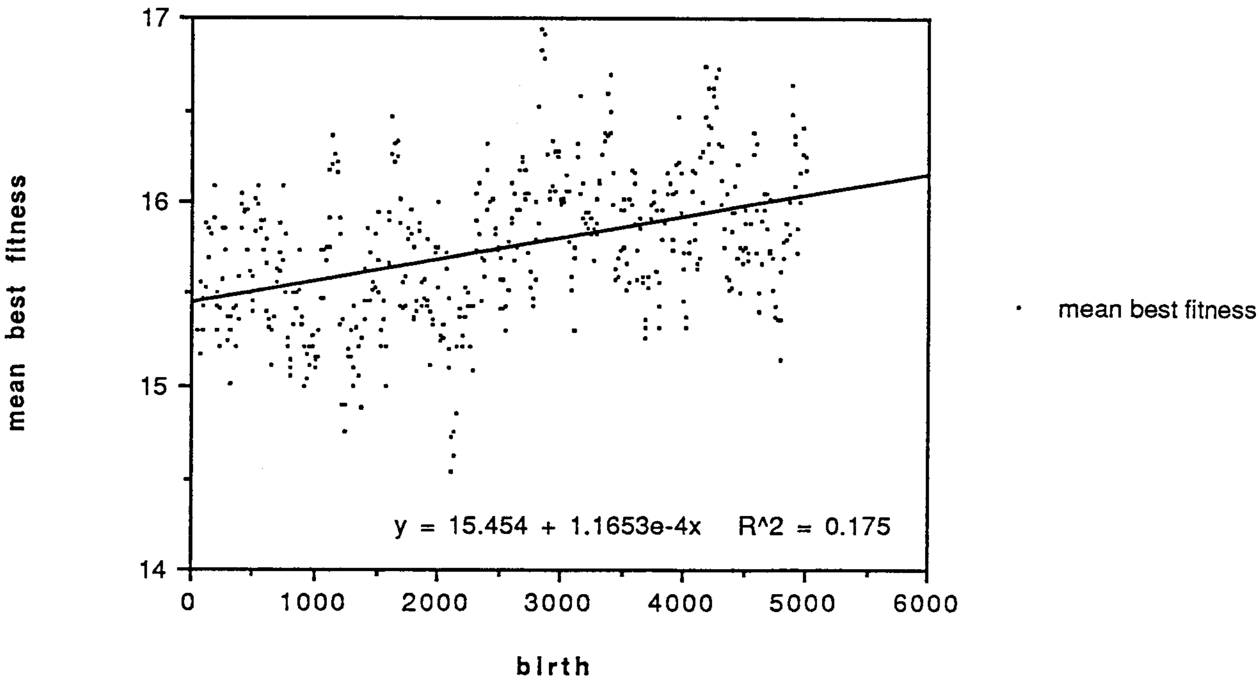

Figure 5: Best Fitness: Communication and Learning Suppressed

(6.25),[10] although linear regression detects a slight upward slope (5.6×10^{-5}). Fitness stays within about 7% of the chance level. The smoothed fitness of the best individual ($\dot{\beta}$) shows a similar lack of direction; see Fig. 5.

Figure 6 shows the evolution of the smoothed average fitness (α) when communication is not suppressed. It begins at the chance level, but by the end of 5000 "weeks" has reached 10.14, which is 62% above the chance level. This is in spite of an apparent "genetic catastrophe" at about $t = 3500$, from which the population did not recover before the termination of the experiment. Linear regression shows that the average fitness is increasing at a rate ($\dot{\alpha}$) that is over 16 times as fast as when communication was suppressed. The smoothed fitness of the best individuals is increasing at an even faster rate ($\dot{\beta}$); see Fig. 7.

Figure 8 shows the evolution of α when both learning and communication are permitted. The rate is three times that when learning was suppressed, and over 50 times that when both communication and learning were suppressed. Figure 9 shows the evolution of β when communication and learning are both permitted.

Finally, Fig. 10 shows the evolution of α when learning is permitted but communication is suppressed. Remarkably, fitness is slowly *decreasing*. We do not yet have an explanation for this phenomenon. Figure 11 shows an analogous decrease in β.

[10]The chance level is that of a random populations; see Section 2.7.1.

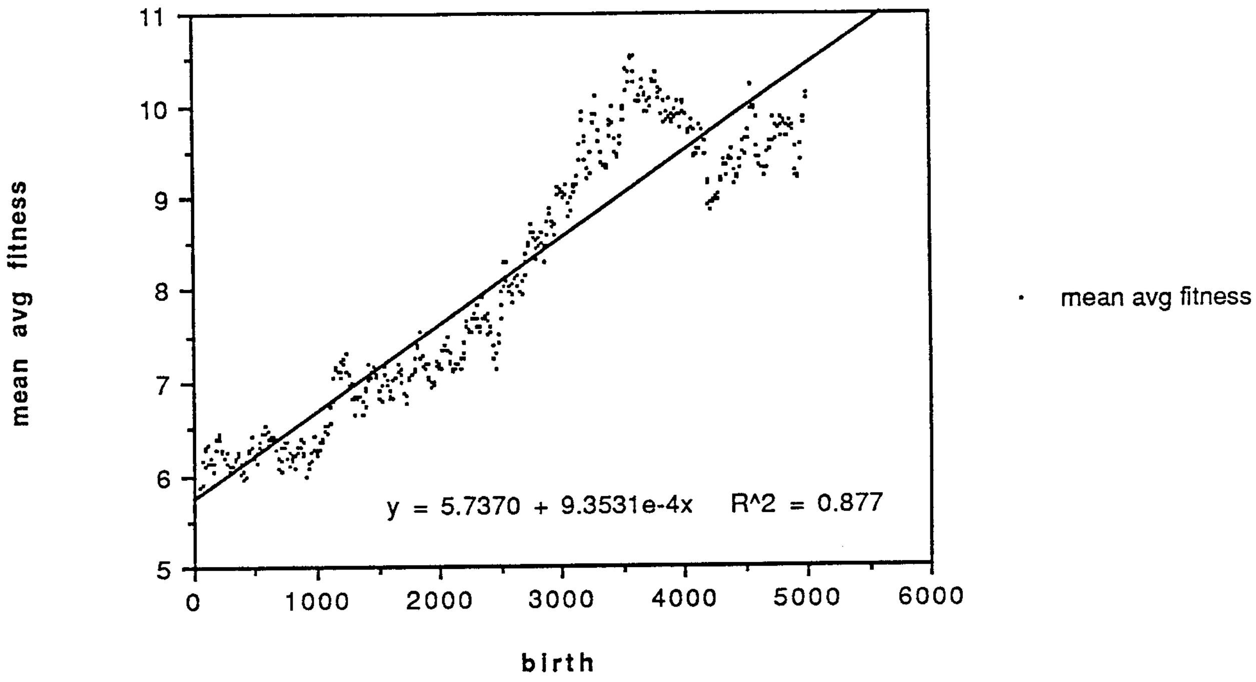

Figure 6: Average Fitness: Communication Permitted, Learning Suppressed

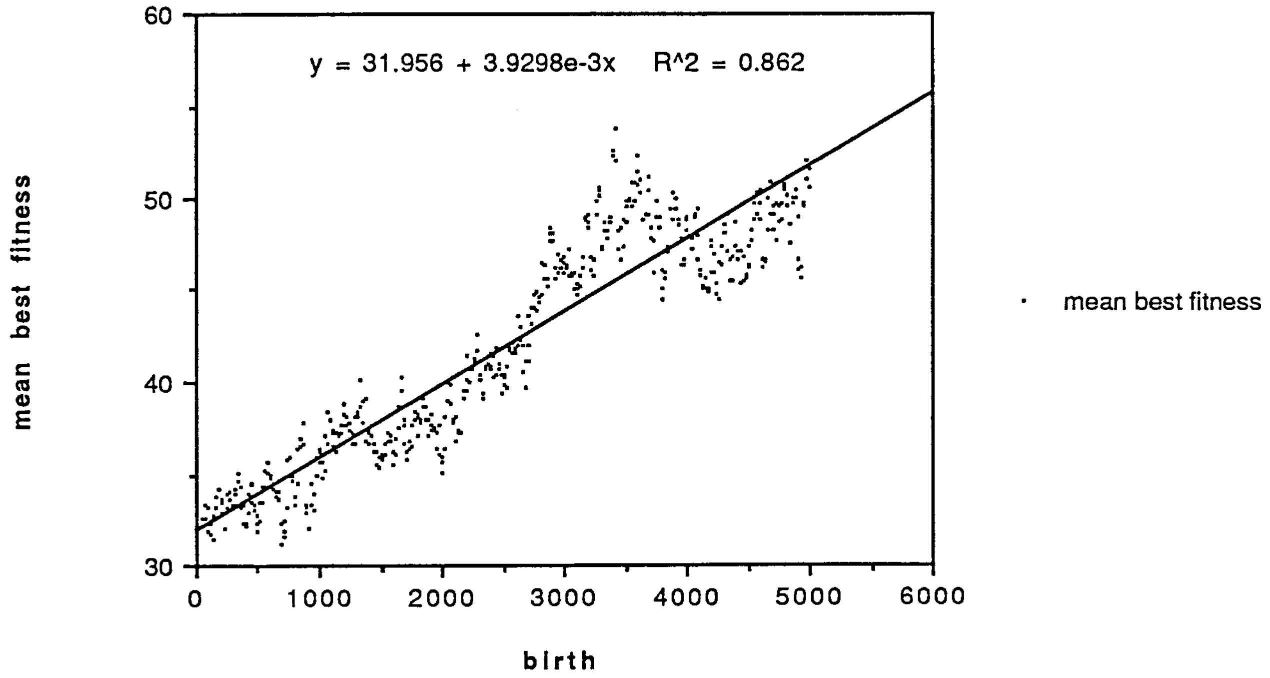

Figure 7: Best Fitness: Communication Permitted, Learning Suppressed

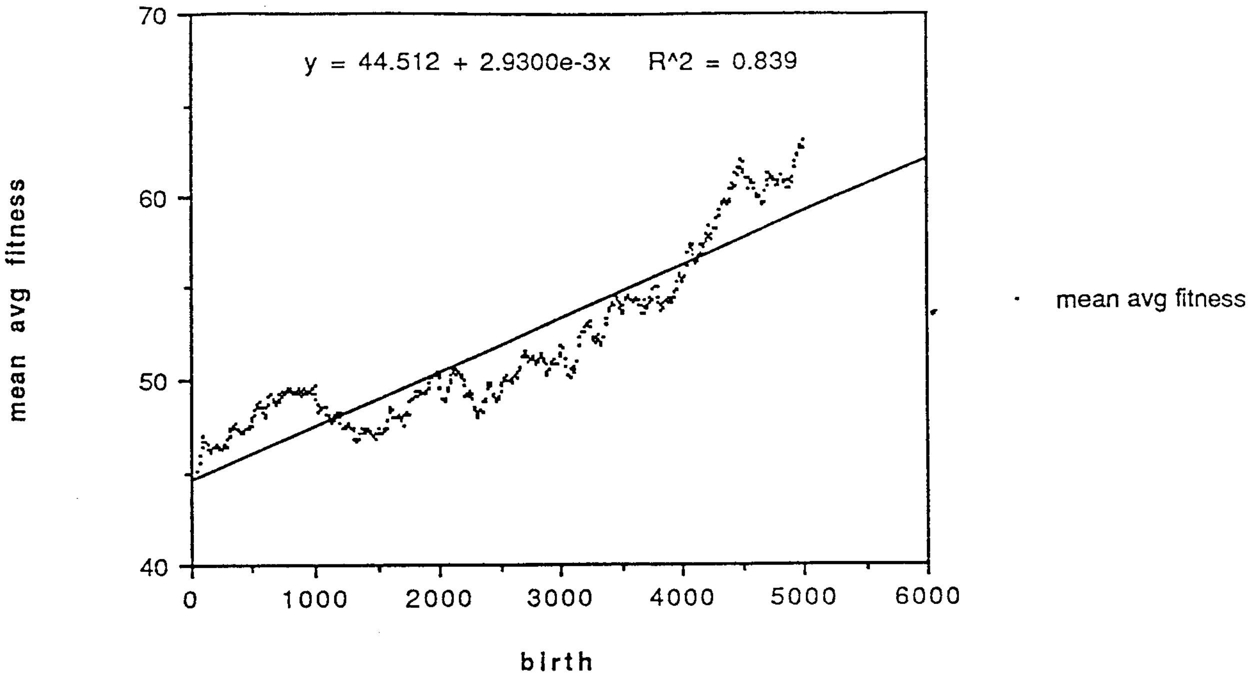

Figure 8: Average Fitness: Communication and Learning Permitted

3.2 Comparison of Denotation Matrices

Next we consider the denotation matrices for the four experiments. Table 3 shows the matrix when communication and learning are both suppressed. It is very uniform, as indicated by its coefficient of variation, $V = 0.52$, and the "unstructuredness parameter" $\eta = 0.92$, which is quite close to 1 (no structure). There is nevertheless a definite pattern in this matrix: each column is quite uniform in value. We have no definitive explanation for this pattern, but expect that it reflects the initial population of random machines.

Table 4 shows the denotation matrix when communication is permitted. Its nonuniformity is apparent to the eye and measured by $V = 2.05$; in other words, the standard deviation is almost twice the mean. Also notice that our measure of lack of structure has decreased to $\eta = 0.35$, which is significantly closer to the "ideal language's" $\eta = 0$.

We can see from the matrix (Table 4) that most of the symbols have fairly specific "meanings," which we can compile into a "dictionary" (Table 5). No symbols are univocal, but a few come close. For example, 90% of the uses of symbol 1 referred to situation 0; all the remaining uses referred to situation 7 (see Table 4). Similarly, 87% of the uses of symbol 0 refer to situation 1. On the other hand, symbol 4 is distinctly equivocal, being used 63% of the time for situation 0 and 37% of the time

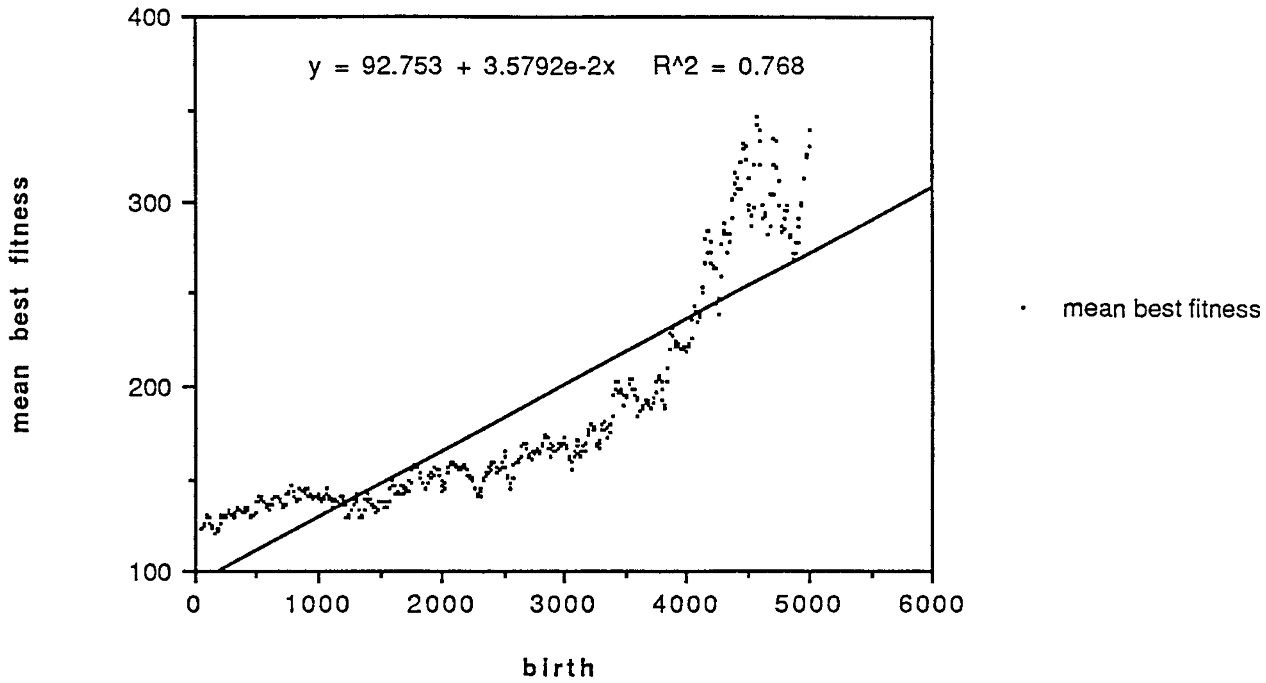

Figure 9: Best Fitness: Communication and Learning Permitted

Table 3: Denotation Matrix: Communication and Learning Suppressed

	situation							
symbol	0	1	2	3	4	5	6	7
0	113	308	129	18	302	349	429	310
1	119	332	145	19	296	357	474	295
2	141	334	145	21	303	342	438	273
3	104	329	165	21	297	362	438	282
4	127	322	150	23	272	327	477	269
5	119	346	167	27	291	341	445	303
6	114	327	145	19	305	372	430	261
7	133	344	166	12	306	321	456	292

$V = 0.527397$
$H = 5.751526$
$\eta = 0.9171753$

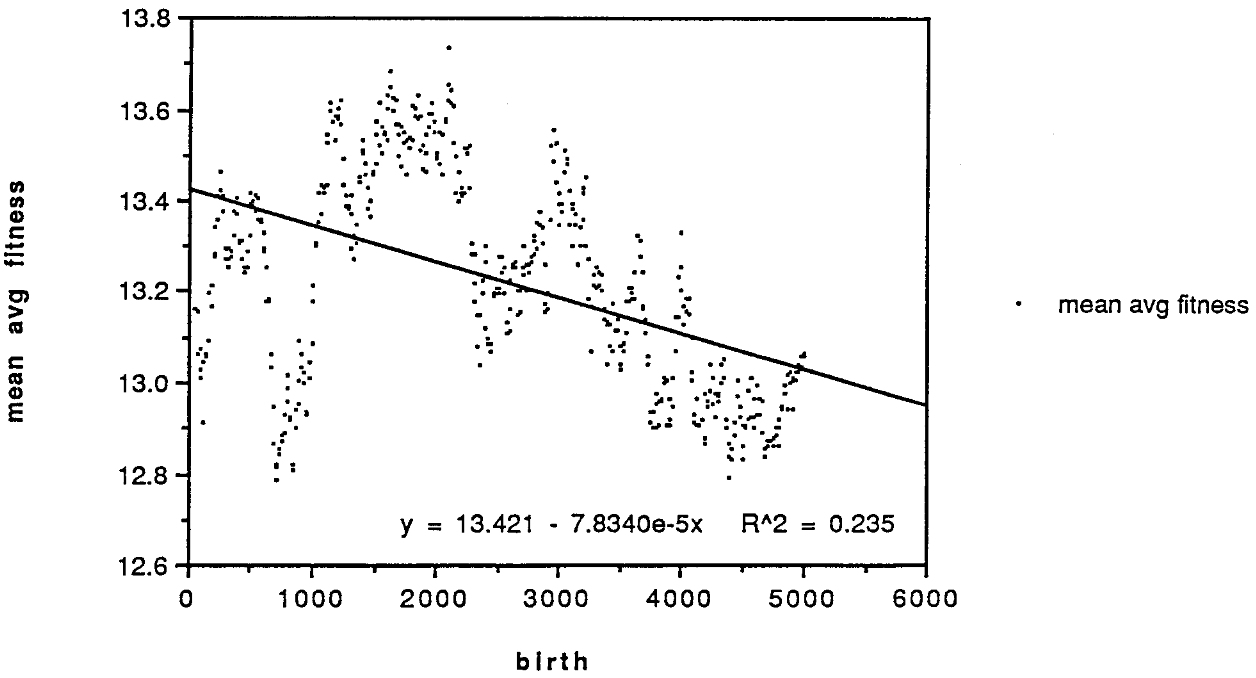

Figure 10: Average Fitness: Communication Suppressed, Learning Permitted

Table 4: Denotation Matrix: Communication Permitted, Learning Suppressed

	situation							
symbol	0	1	2	3	4	5	6	7
0	0	2726	0	0	168	130	82	30
1	896	0	0	0	0	0	0	100
2	0	0	278	1101	265	99	0	0
3	270	0	0	611	4	0	4192	39
4	2530	0	0	0	1	0	0	1492
5	0	397	806	1896	1	0	10	210
6	0	70	371	619	2921	0	0	0
7	935	202	0	498	0	0	1194	212

$V = 2.054575$
$H = 4.063197$
$\eta = 0.354399$

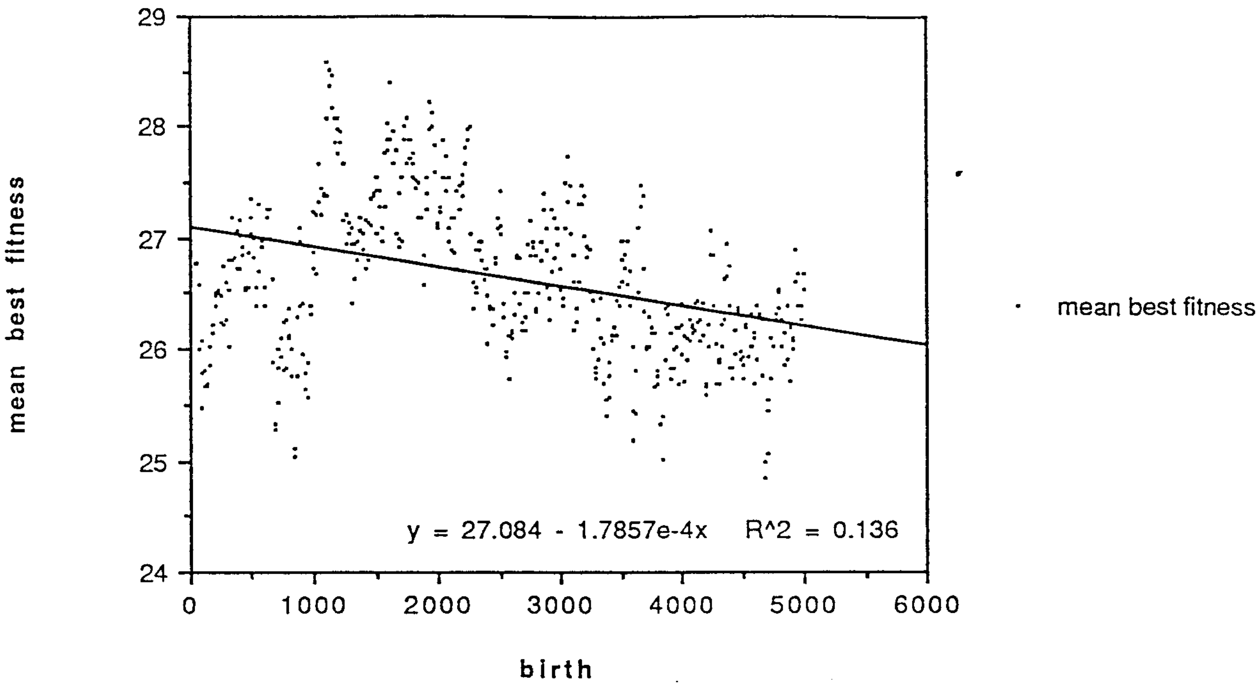

Figure 11: Best Fitness: Communication Suppressed, Learning Permitted

Table 5: Dictionary Compiled From Denotation Matrix

symbol	situation
0	1
1	0
2	3
3	6
4	0 or 7
5	3
6	4
7	0 or 6

Table 6: Denotation Matrix: Communication and Learning Permitted

	situation							
symbol	0	1	2	3	4	5	6	7
0	27	2845	393	0	2258	179	0	0
1	0	516	0	3864	1221	987	1662	4532
2	296	2756	198	844	1893	0	1374	0
3	0	817	0	216	0	0	2193	126
4	36	3365	7926	547	0	1143	444	0
5	0	0	100	297	0	103	63	778
6	1	3936	31095	24580	4780	28302	2086	1589
7	3685	2	0	2	768	2603	4273	5762

$$V = 2.411387$$
$$H = 4.011914$$
$$\eta = 0.3373047$$

for situation 7.

Note that this observed ambiguity may reflect either one "language community" using the symbol for two different situations, or two communities, each with its own language. We cannot distinguish these possibilities from the denotation matrix alone. Instead, we must "dissect" the actual machines constituting the final population, a process demonstrated below (Section 3.3).

Finally, observe that there are no symbols that refer unambiguously to situations 2 or 5, although there is a symbol (viz. 5) that often refers to situation 2. There are no symbols in this language for situation 5.

Table 6 shows the denotation matrix that resulted when both communication and learning were permitted. It is even more nonuniform than Table 4, with $V = 2.41$ and $\eta = 0.34$. On the other hand, it is somewhat harder to extract a dictionary from this matrix, perhaps reflecting language instability that could result from learning. In other words, the ability to learn permits *pro tempore* languages to be set up that will function adequately for a "day" (i.e. until the next environment change).

Table 7 shows the denotation matrix that resulted from suppressing communication but permitting learning. As expected, it is very uniform ($V = 0.53$, $\eta = 0.92$).

For comparison, Figures 12, 13, 14 and 15 show average fitness evolution for a different random population. The denotation matrices are in Tables 8, 9, 10 and 11. They do not differ qualitatively from what we've seen.

From the denotation matrix in Table 9 we can extract two dictionaries, the *reception dictionary* (Table 12), which maps symbols into situations, and the *emission dictionary* (Table 13), which maps situations into symbols. In each case we pick the largest entry on the row (for a symbol) or column (for a situation), unless the largest

Table 7: Denotation Matrix: Communication Suppressed, Learning Permitted

	situation							
symbol	0	1	2	3	4	5	6	7
0	1058	322	564	277	692	641	74	429
1	1036	316	517	267	703	654	89	425
2	1066	299	547	324	681	636	88	413
3	1085	335	628	320	651	647	94	403
4	1019	324	570	302	668	676	85	422
5	1019	342	577	295	673	644	78	427
6	1091	334	556	287	628	619	71	431
7	1125	315	589	295	692	697	90	435

$V = 0.549618$
$H = 5.769859$
$\eta = 0.9232863$

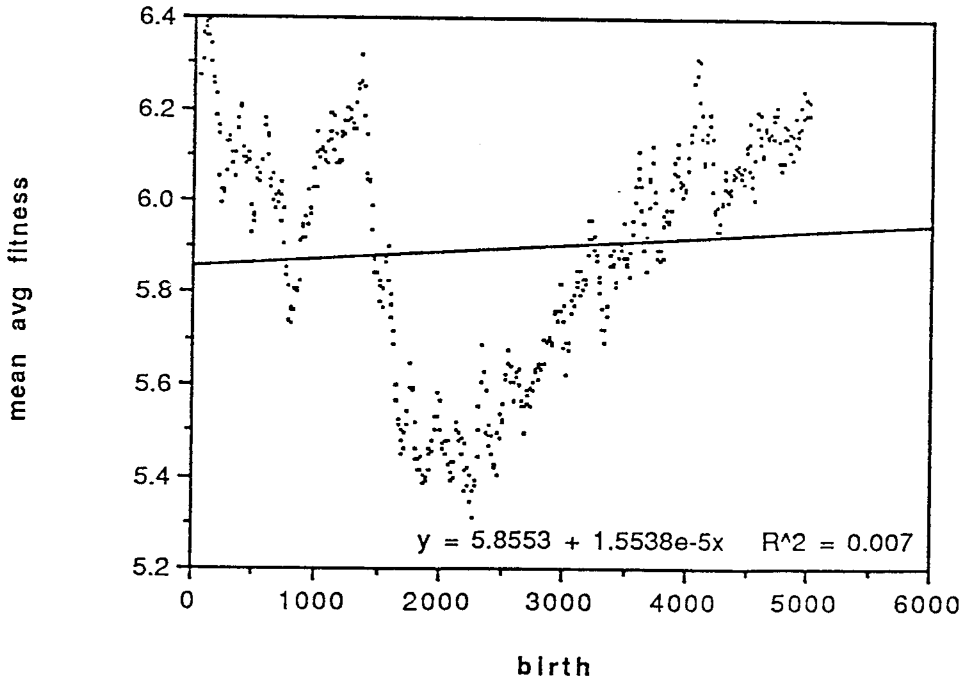

Figure 12: Average Fitness: Communication and Learning Suppressed

Table 8: Denotation Matrix: Communication and Learning Suppressed

	situation							
symbol	0	1	2	3	4	5	6	7
0	320	138	189	360	266	354	224	89
1	364	130	189	359	261	342	266	75
2	332	126	184	385	252	365	257	82
3	350	125	193	366	257	351	255	98
4	340	119	190	354	254	356	225	78
5	328	145	170	343	244	348	217	86
6	345	119	194	374	214	361	237	78
7	346	149	159	343	242	383	226	83

$V = 0.409451$
$H = 5.868233$
$\eta = 0.9560777$

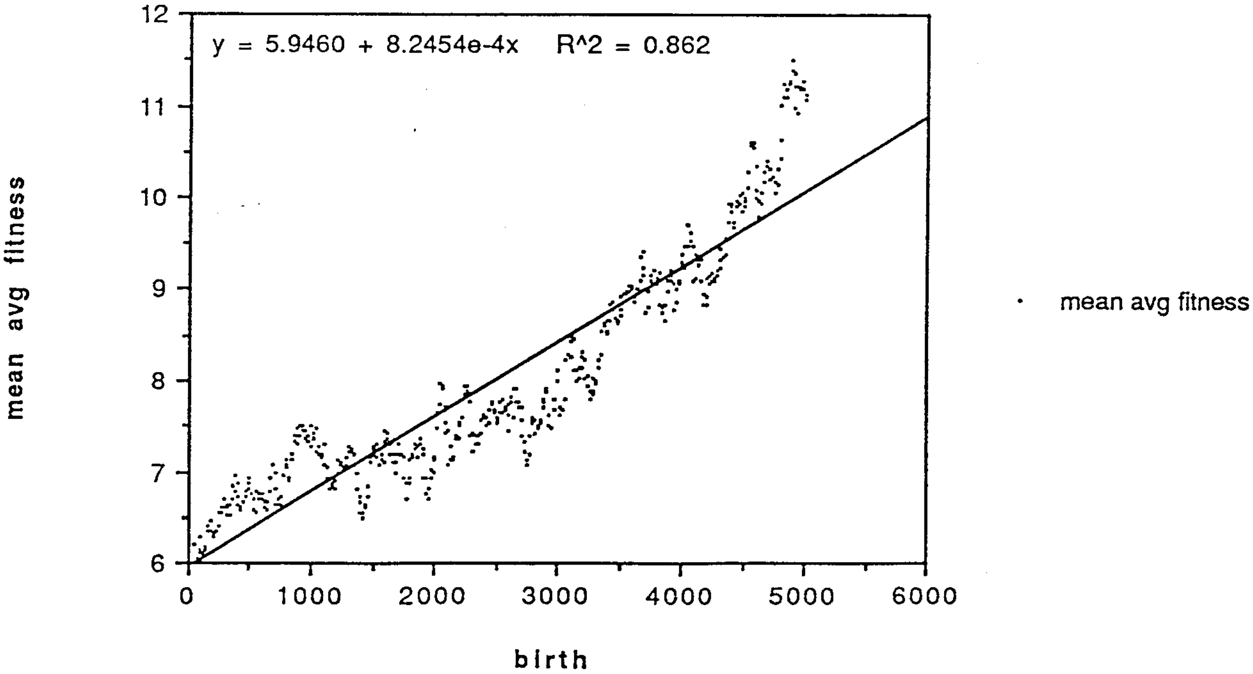

Figure 13: Average Fitness: Communication Permitted, Learning Suppressed

Table 9: Denotation Matrix: Communication Permitted, Learning Suppressed

	situation							
symbol	0	1	2	3	4	5	6	7
0	695	5749	0	1157	0	2054	101	0
1	4242	11	1702	0	0	0	1	0
2	855	0	0	0	0	603	862	20
3	0	0	0	0	1003	430	0	1091
4	0	0	0	0	0	0	2756	464
5	0	0	40	0	548	0	817	0
6	1089	90	1	281	346	268	0	62
7	0	201	0	288	0	0	2	0

$V = 2.272352$
$H = 3.915812$
$\eta = 0.3052707$

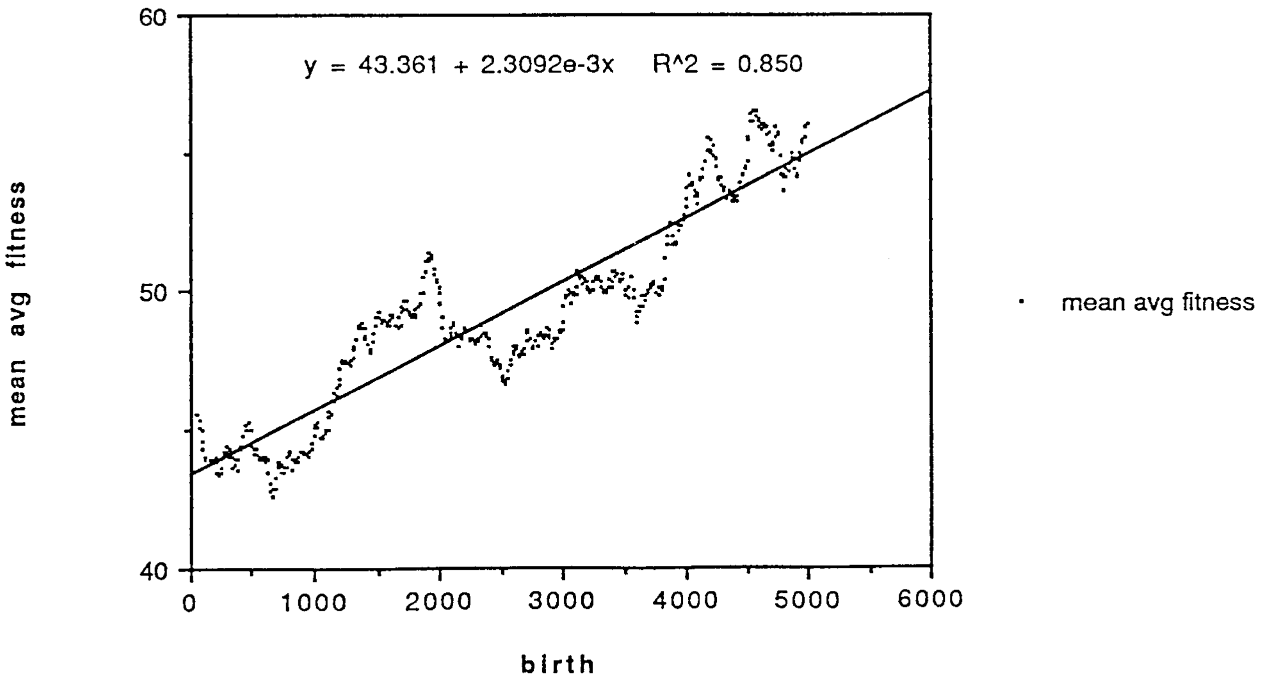

Figure 14: Average Fitness: Communication and Learning Permitted

Table 10: Denotation Matrix: Communication and Learning Permitted

	situation							
symbol	0	1	2	3	4	5	6	7
0	0	0	2946	0	0	635	4239	3233
1	2084	0	672	1457	0	6701	8517	1284
2	0	0	646	433	0	230	63	879
3	0	1074	446	46	2315	1623	0	1265
4	27850	5504	0	2326	11651	243	3428	20076
5	1301	0	0	854	858	368	0	0
6	13519	2676	0	2223	2391	874	0	644
7	356	226	365	107	1357	27	100	1

$V = 2.165397$
$H = 4.208782$
$\eta = 0.4029273$

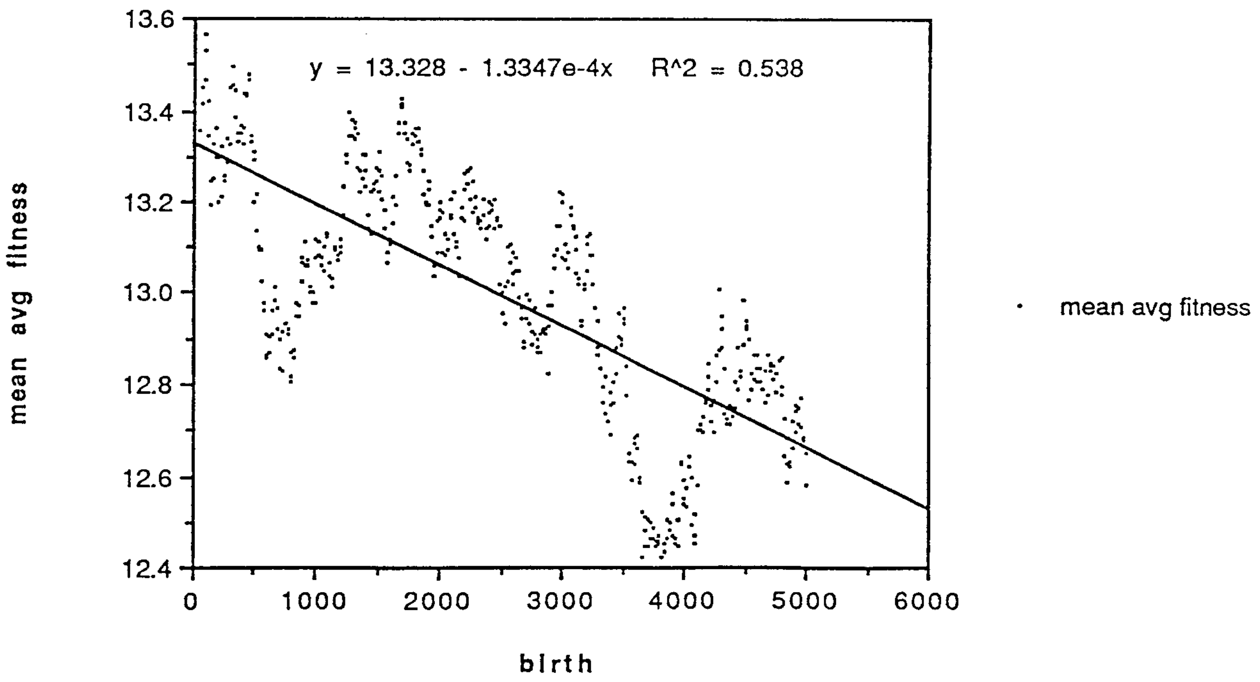

Figure 15: Average Fitness: Communication Suppressed, Learning Permitted

Table 11: Denotation Matrix: Communication Suppressed, Learning Permitted

	situation							
symbol	0	1	2	3	4	5	6	7
0	173	429	398	591	358	596	745	594
1	152	406	444	612	386	560	752	559
2	148	444	461	668	371	590	700	565
3	174	441	476	636	367	552	736	566
4	178	448	464	641	351	554	757	522
5	160	464	437	665	350	567	705	582
6	168	486	463	630	387	570	762	572
7	183	474	438	620	383	539	724	558

$$V = 0.335501$$
$$H = 5.90837$$
$$\eta = 0.9694567$$

Table 12: Reception Dictionary

symbol	$\longmapsto$	situation
0		1
1		0
2		0 or 6
3		4 or 7
4		6
5		6
6		0
7		1 or 3

Table 13: Emission Dictionary

situation	$\longmapsto$	symbol
0		1
1		0
2		1
3		0
4		3
5		0
6		4
7		3

Table 14: Bidirectional Dictionary

symbol	$\longleftrightarrow$	situation
0		1
1		0
3		4 or 7
4		6

two are nearly equal, in which case we list both.[11]

Notice that the reception and emission dictionaries are not inverses of each other; there is no requirement that a symbol mean the same thing to an organism when it is received as when it is emitted. However, by comparing Tables 12 and 13 we can see that half of the symbols are in fact used bidirectionally; these are shown in Table 14.

Observe that this population has evolved a language in which a symbol ambiguously denotes situations 4 and 7, *in either direction.* This is remarkable consistency of usage.

3.3 Analysis of Phenotypes

If Tables 9, 12, 13 and 14 reflect the language being used by this population, then we ought to be able to see some evidence of it in the structure of the machines. Therefore we have "dissected" the most fit individual from this experiment (Experiment ID #1213).

In Table 15 we see this machine's transition table listed by input symbol (global environment state); this is the machine's effective reception dictionary. Table 16 lists the same transition table by situation (local environment state); this is its effective emission dictionary. Each triple represents (1) the new internal state (always 0 in these experiments), (2) the kind of response (0 = act, 1 = emit), and (3) the action or emitted symbol.

We can now compare the reception dictionary extracted from the population (Table 12) with that implicit in the most fit individual (Table 15). There are 13 matches; 5.5 would be expected by chance.[12] (The matches are shown in boldface.) Similarly we can compare the population and individual emission dictionaries; there are 12 matches where only 4 would be expected (Table 16).

Obviously this kind of "dissection" of individual machines is quite laborious. Automatic tools need to be developed for analyzing the evolved structure and for gathering

[11] Notice that Table 13 shows that symbol 0 almost means "lowest bit is 1" since we have $1 \mapsto 0$, $3 \mapsto 0$ and $5 \mapsto 0$.

[12] 11 entries on the right of Table 12 $\times$ 1/2 expected match each (8 chances to hit 1 of 16 possibilities).

Table 15: Reception Dictionary of Best Individual

	situation							
symbol	0	1	2	3	4	5	6	7
0	0 1 1	0 1 0	0 0 3	0 0 5	**0 0 1**	0 0 0	0 1 4	**0 0 1**
1	0 0 2	0 1 6	0 1 5	0 1 7	0 1 3	0 1 0	**0 0 0**	**0 0 0**
2	**0 0 0**	0 1 0	0 1 2	0 1 3	0 1 3	0 0 5	**0 0 6**	0 1 3
3	0 1 2	**0 0 7**	0 1 1	0 1 0	0 0 5	0 1 3	0 1 5	**0 0 4**
4	0 1 2	0 0 5	0 1 2	0 1 6	**0 0 6**	0 1 6	0 1 4	0 0 3
5	0 1 0	0 0 4	0 1 1	0 1 6	**0 0 6**	0 1 7	**0 0 6**	0 1 5
6	0 0 4	0 0 5	0 1 7	0 1 3	0 0 7	0 1 2	**0 0 0**	0 1 1
7	0 1 6	0 1 2	**0 0 1**	0 1 6	0 1 2	0 1 6	0 1 0	0 1 2

Table 16: Emission Dictionary of Best Individual

	symbol							
situation	0	1	2	3	4	5	6	7
0	**0 1 1**	0 0 2	0 0 0	0 1 2	0 1 2	0 1 0	0 0 4	0 1 6
1	**0 1 0**	0 1 6	**0 1 0**	0 0 7	0 0 5	0 0 4	0 0 5	0 1 2
2	0 0 3	0 1 5	0 1 2	**0 1 1**	0 1 2	**0 1 1**	0 1 7	0 0 1
3	0 0 5	0 1 7	0 1 3	**0 1 0**	0 1 6	0 1 6	0 1 3	0 1 6
4	0 0 1	**0 1 3**	**0 1 3**	0 0 5	0 0 6	0 0 6	0 0 7	0 1 2
5	0 0 0	**0 1 0**	0 0 5	0 1 3	0 1 6	0 1 7	0 1 2	0 1 6
6	**0 1 4**	0 0 0	0 0 6	0 1 5	**0 1 4**	0 0 6	0 0 0	0 1 0
7	0 0 1	0 0 0	**0 1 3**	0 0 4	0 0 3	0 1 5	0 1 1	0 1 2

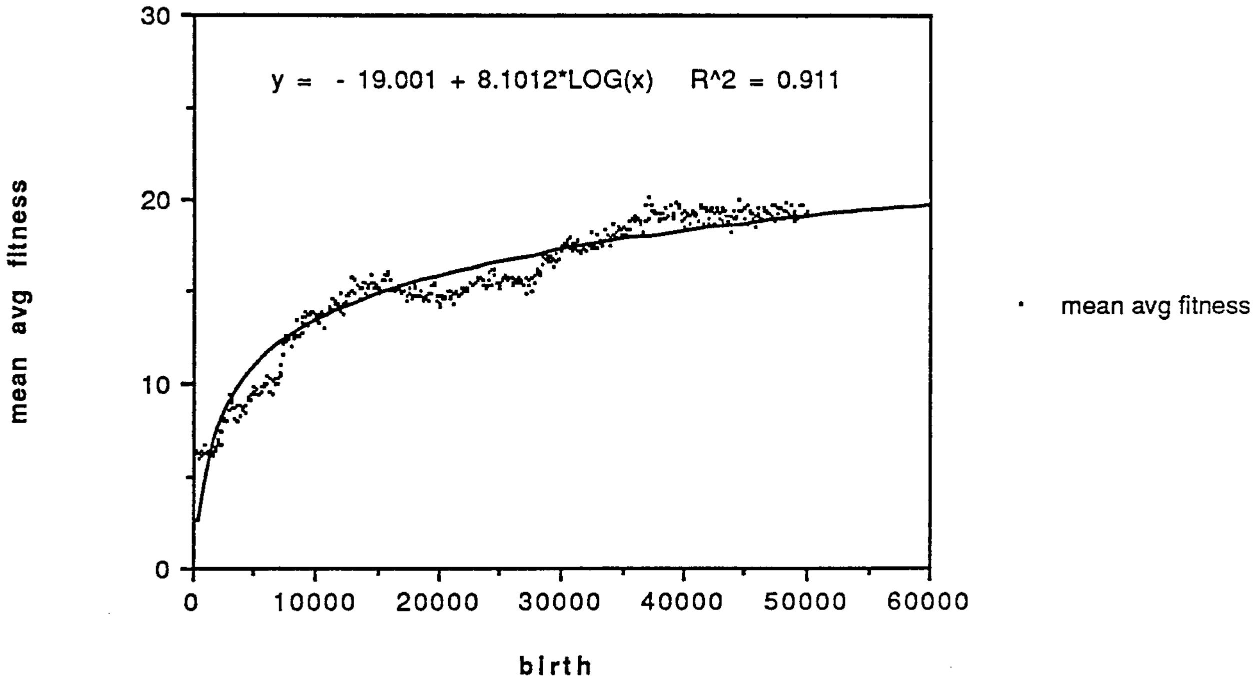

Figure 16: Average Fitness: Communication Permitted, Learning Suppressed

more sophisticated statistics.

3.4 Additional Experiments

Figures 16 – 19 show the fitness evolution for the longest simulations run to date: 50000 (5×10^4) "weeks." In all of these it is apparent that the rate is decreasing, so it seems reasonable to fit a log curve to the data. Notice especially that in Figs. 18 and 19 the fitness seemed to have reached a plateau at about $t = 15000$, which continued until the "catastrophe" at about $t = 45000$. Presumably the population would have recovered from this had the experiment continued. In any case, it appears that $\alpha = 57$ and $\beta = 210$ are the equilibrium values.[13] The final denotation matrices are in Tables 17 and 18. From Table 17 we can see that symbols 1, 2, 3 and 5 have distinct meanings, symbol 4 has two clear meanings, and symbols 0, 6 and 7 are not used at all. This is reflected in the value $\eta = -0.2$, which indicates an "over-structured" language (see p. 12). The lower coefficient of variation when learning was permitted ($V = 2.23$, Table 18) was presumably a result of the "catastrophe."

Table 19 lists final fitness (α, β), rate of fitness change ($\dot{\alpha}$, $\dot{\beta}$) and measure of

[13]More specifically, the averages of the values from $t = 15000$ to $t = 45000$ are $\bar{\alpha} = 56.634$ and $\bar{\beta} = 210.380$.

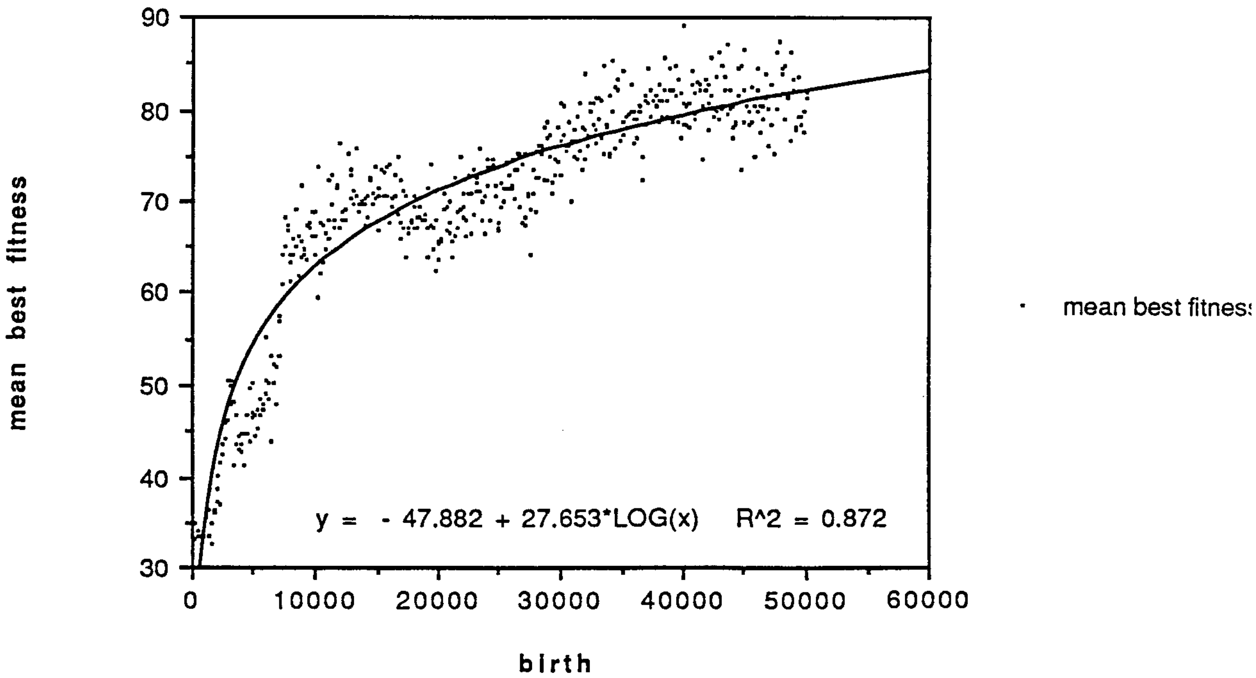

Figure 17: Best Fitness: Communication Permitted, Learning Suppressed

Table 17: Denotation Matrix: Communication Permitted, Learning Suppressed

	situation							
symbol	0	1	2	3	4	5	6	7
0	0	0	0	0	0	0	0	0
1	0	0	11014	5	1393	0	2	0
2	9	3	0	0	0	0	0	2494
3	0	2	0	0	0	629	0	1
4	0	181	0	0	3560	0	4696	1
5	0	0	0	20334	0	4	0	3898
6	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0

$V = 3.914863$
$H = 2.398098$
$\eta = -0.200634$

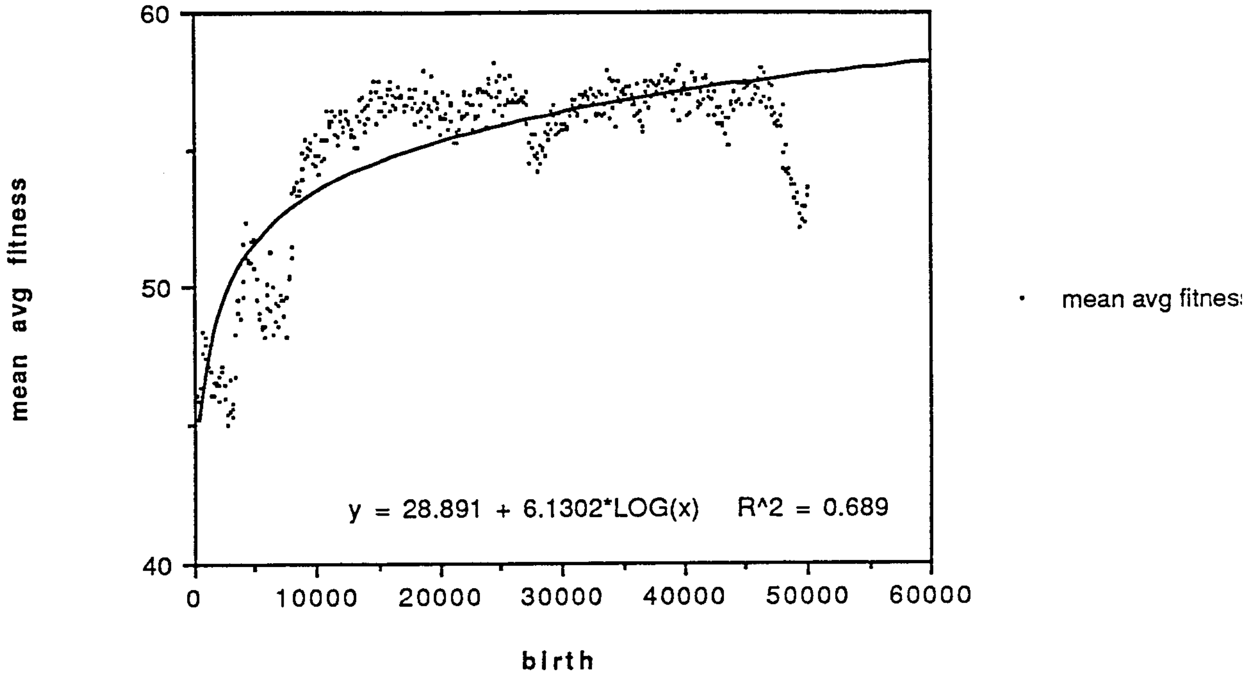

Figure 18: Average Fitness: Communication and Learning Permitted

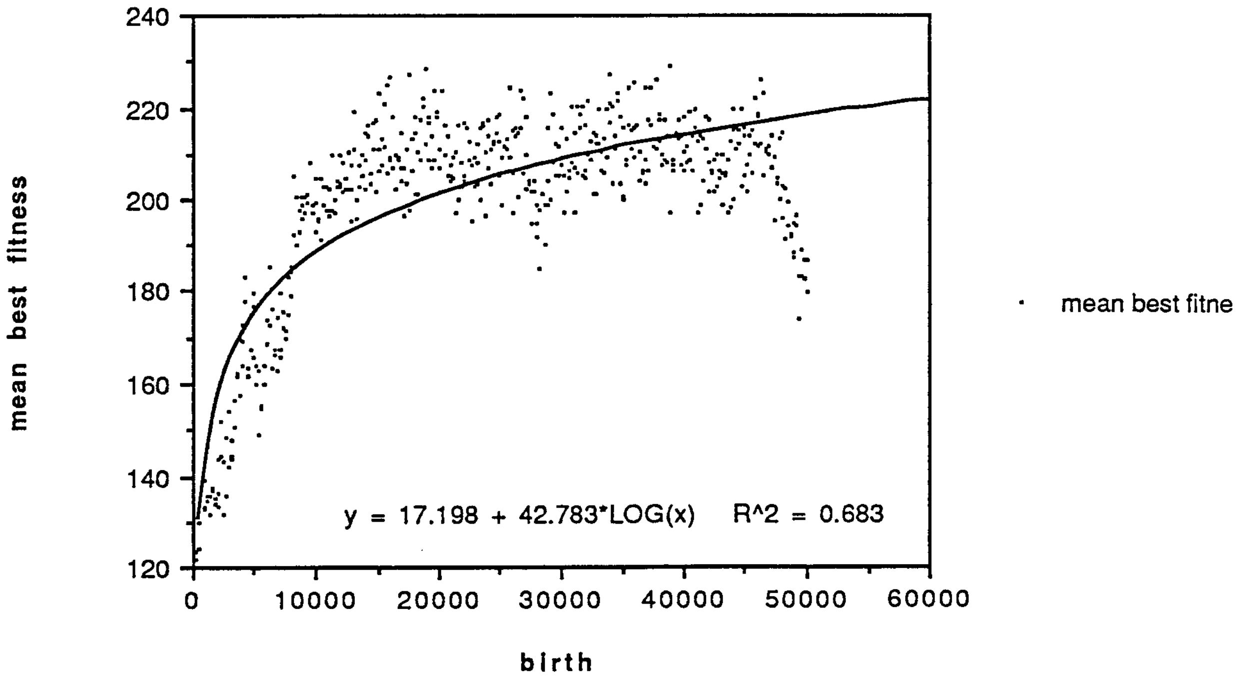

Figure 19: Best Fitness: Communication and Learning Permitted

Table 18: Denotation Matrix: Communication and Learning Permitted

	situation							
symbol	0	1	2	3	4	5	6	7
0	0	2	0	8907	0	0	0	0
1	0	0	1070	0	606	0	4667	0
2	0	16647	21703	0	1	0	0	0
3	0	1666	2440	0	0	1358	0	882
4	4695	1339	1	7445	330	1277	0	0
5	660	0	0	0	0	0	0	90
6	0	0	1	16793	270	12149	0	21167
7	0	0	0	2121	2785	0	337	2563

$V = 2.358346$
$H = 3.685994$
$\eta = 0.2286647$

Table 19: Summary of Fitness Rates for Individual Experiments

ID	Comm	Learn	α	β	$\dot{\alpha}$	$\dot{\beta}$	V	H	η
1217a	N	N	6.22	15.82	0.16	0.53	0.41	5.87	0.96
0104a	N	N	6.40	16.18	0.56	1.17	0.53	5.75	0.92
1213	Y	N	11.13	53.74	8.2	31.4	2.27	3.92	0.31
1214	Y	N	12.78	65.22	13.0	72.4	4.00	2.93	−0.02
1215	Y	N	12.48	57.70	13.4	49.2	2.01	4.21	0.40
0104	Y	N	10.14	50.52	9.4	39.3	2.05	4.06	0.35
1215b	Y	Y	56.05	223.28	23.1	199.7	2.17	4.21	0.40
1216	Y	Y	59.90	219.72	33.9	213.3	3.37	3.39	0.13
0104b	Y	Y	62.99	339.84	29.3	357.9	2.41	4.01	0.34
0104c	N	Y	13.06	26.68	−0.78	−1.79	0.55	5.77	0.92
0107	N	Y	12.58	24.44	−1.33	−3.42	0.34	5.91	0.97

Notes:
ID = experiment identification number
Comm = communication permitted
Learn = learning permitted
$\dot{\alpha}$ and $\dot{\beta}$ scaled by 10^4

Table 20: Summary of Fitness Evolution

	Comm/Learn			
	N/N	Y/N	Y/Y	N/Y
α	6.31	11.63	59.65	12.82
β	16.00	56.80	260.95	25.56
$\dot{\alpha}$	0.36	11.0	28.77	−1.06
$\dot{\beta}$	0.85	48.1	257.0	−2.61
V	0.47	2.58	2.65	0.44
H	5.81	3.79	3.87	5.84
η	0.94	0.26	0.29	0.95
α ratio	1	1.84	9.45	2.03
β ratio	1	3.55	16.31	1.60
$\dot{\alpha}$ ratio	1	30.6	79.9	−2.94
$\dot{\beta}$ ratio	1	56.6	302.4	−3.07
V ratio	1	5.52	5.66	0.94
H^{-1} ratio	1	1.53	1.50	0.99
η^{-1} ratio	1	3.62	3.24	0.99

Note: $\dot{\alpha}$, $\dot{\beta}$ scaled by 10^4

structure (V, H, η) for several experiments that were run for 5000 "weeks." The table indicates whether communication and learning were permitted. In Table 20 we have averaged the runs with the same communication and learning parameters so that their effect is more apparent. In the remainder of this section we summarize these effects.

With communication permitted, the average fitness increases at a rate over 30 times as fast as when it is not. Also, the fitness of the best individuals increases as a rate over 50 times as fast, the coefficient of variation is over five times as large, the entropy is significantly decreased (by a factor of 1.5), and the structure measure (η^{-1}) is over three times as large.

When both communication and learning are permitted, average fitness increases at a rate 80 times as fast as when neither is permitted. Best fitness increases at a rate over 300 times as fast, but the coefficient of variation is about the same as with no learning and entropy is slightly higher (perhaps reflecting more rapid linguistic change).

On the other hand, when learning was permitted but communication suppressed, the fitness actually *decreased*, albeit slowly. At this time we do not have an explanation for this phenomenon, nor indeed for the slow increase of fitness when both communication and learning are suppressed. In both cases there should be no selective pressure, since it is impossible to improve on guessing. However, it must be pointed out that the rates are very slow. For example, in the case where communication and learning are both suppressed, the average fitness stays very close to the chance level (within 1%),[14] so perhaps the change is a result of genetic drift.

As expected, when communication is suppressed, the suppression or not of learning has little effect on the structure of the language; V, H and η are all similar in the two cases.

4 Future Investigations

There are many issues that need further investigation. Some of these can be settled by gathering additional statistics. For example, we would like to know the fraction of emissions that lead to successful communication acts. We expect this to be $1/L$ when communication is suppressed (due to guessing), but significantly higher when communication is permitted.[15] It would also be interesting to compare the number of communication acts that result from learning as opposed to inheritance.

We want to investigate several kinds of learning. For example, the current "single trial learning" can lead to instability in the phenotype. Perhaps it would be preferable to require some reinforcement before learning takes place. Also, machines now learn through their own trial and error, but it would be simple to have them learn from other, more fit machines, thus permitting true cultural inheritance [3].

"Dissecting" individual machines is labor intensive, so we would like to be able to automate the process. We would also like more systematic comparisons of the structure of the machines and the statistical distribution of communication acts. For example, we would expect that the denotation matrix implicit in the population (perhaps weighted by fitness) should be highly correlated to the denotation matrix derived from the simulation.

In all the experiments described in this report we have taken $G = L$, that is, the number of possible symbols is the same as the number of possible situations. Suppose we take $G > L$; then there are more symbols than we need. Will we find some symbols being unused? Or will we find synonymous symbols? Or multiple language communities using different symbols for the same situation?

On the other hand suppose we take $G < L$; then there are too few symbols. In this case, if the machines have more than one internal state, we might find that the population begins to string symbols together to denote situations. What syntax will they use? Will "word order" be significant? Will there be a definite grammar? Will more than one language evolve? If so, how will they interact?

[14] The average α reached for the two experiments was 6.31. Since the chance level is 6.25, we have $6.31/6.25 = 1.0096$.

[15] In fact this ratio is currently computed, but over the entire simulation, not just the last W "weeks." Therefore, the ratio is dominated by the early phases of the evolution, before communication has emerged.

If environmental situations had features that were somewhat independent, would separate symbols for these features emerge? Would there be any evidence of linguistic categories (e.g. nouns, verbs, modifiers)?

Suppose we impose a spatial metric on the environments (so that some are closer than others), and make probability of communicating and breeding decrease with distance. Will we find geographically local languages evolving? What will be the dynamics of their boundaries?

We anticipate a number of interesting differences would result from using neural networks, rather than finite state machines, to determine the behavior of the individuals in the population.[16] The sources of these differences include the continuous response of neurons, which may result in "continuous" languages, and the ability to use the genotype to govern a more complicated developmental process than has been the case so far. Neural networks would also permit more realistic investigations of learning.

5 Conclusions

We have shown that communication may evolve in a population of simple machines that are physically capable of sensing and modifying a shared environment, and for which there is selective pressure on cooperative behavior. The emergence of communication was detected by comparing simulations in which communication was permitted with those in which it was suppressed. When communication was not suppressed we found that at the end of the experiment the average fitness of the population was 84% higher and had increased at a rate 30 times faster than when communication was suppressed. Furthermore, when communication was suppressed, the statistical association of symbols with situations was random, as was expected. In contrast, permitting communication led to very structured associations of symbols and situations, as determined by a variety of measures (V, H, η). Inspection of the structure of individual highly fit machines confirmed the statistical structure.

Our simulations also investigated a simple kind of learning. This did not help (and in fact hurt) when communication was suppressed, but when communication was permitted the resulting fitness was 845% higher and increased at a rate 80 times as fast as when it was suppressed.

Finally, we believe that the experiments described here show a new way to investigate the emergence of communication, its function in populations of simple machines, and the structure of the resulting symbol systems.

[16] In fact this was our original intent; using FSMs was intended as a preliminary investigation to develop the required procedures.

References

[1] Baldwin, J. M., A New Factor in Evolution, *American Naturalist* **30** (1896), 441-451.

[2] Belew, Richard K., Evolution, Learning and Culture: Computational Metaphors for Adaptive Algorithms, University of California, San Diego, Computer Science & Engineering Dept., #CS89-156 (September 1989).

[3] Boyd, Richard, and Richerson, Peter J., *Culture and the Evolutionary Process*, University of Chicago Press, Chicago, 1985.

[4] Burghardt, Gordon M., Defining "Communication," *Communication by Chemical Signals*, J. W. Johnston Jr., D. G. Moulton and A. Turk (eds.), Appleton-Century-Crofts, New York, 1970.

[5] Dawkins, Richard, *The Extended Phenotype*, Oxford University Press, Oxford, 1982.

[6] Goldberg, D. E., *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, 1989.

[7] Hinton, Geoffrey E., and Nowlan, Steven J., How Learning Can Guide Evolution, *Complex Systems* **1** (1987), 495-502.

[8] Holland, J. H., *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, 1975.

[9] MacLennan, Bruce J., "Causes and Intentions," *Behavioral and Brain Sciences*, **11**, 3 (September 1988), 519-520.

[10] Maynard-Smith, John, When Learning Guides Evolution, *Nature* **329** (Oct. 29 1987), 761-762.