

DRAFT

Not to be Reproduced or Generally Distributed

Submitted to *Adaptive Behavior*

Synthetic Ethology and the Evolution of Cooperative Communication

Bruce J. MacLennan* & Gordon M. Burghardt†

*Computer Science Department

University of Tennessee, Knoxville TN 37996, U.S.A

MacLennan@CS.UTK.EDU

(615)974-5067

FAX: (615)974-4404

†Psychology Department & Graduate Program in Ethology

University of Tennessee, Knoxville TN 37996, U.S.A

BURGH@UTKVX.UTK.EDU

Abstract. *Synthetic ethology* is proposed as a means of conducting controlled experiments investigating the mechanisms and evolution of communication. After a discussion of the goals and methods of synthetic ethology, two series of experiments are described based on at least 5000 breeding cycles. The first demonstrates the evolution of cooperative communication in a population of simple machines. The average fitness of the population and the organization of its use of signals are compared under three conditions: communication suppressed, communication permitted, and communication permitted in the presence of learning. Where communication is permitted the fitness increases about 26 times faster than when communication is suppressed; with communication and learning the rate of fitness increase is about 100 fold. The second series of experiments illustrates the evolution of a syntactically simple language, in which a pair of signals is required for effective communication.

Keywords: artificial life, communication, cooperation, entropy, ethology, evolution, genetic algorithm, intentionality, language, learning, synthetic ethology

Running Title: Synthetic Ethology

Introduction

The Role and Scope of Synthetic Ethology

Description of Synthetic Ethology

Synthetic ethology is an approach to the study of animal behavior in which simple, synthetic organisms are allowed to behave and evolve in a synthetic world. Because both the organisms and their worlds are synthetic, they can be constructed for specific purposes, in particular, for testing specific hypotheses. This approach permits more carefully controlled experiments than are otherwise possible in ethology, because variables can be controlled more precisely, and the evolution of identical populations can be studied under differing conditions. In this paper we demonstrate the application of synthetic ethology methods to the evolution of cooperative communication. Through the analysis of this important topic in theoretical and empirical ethology, we want to stimulate discussion of the general utility of the synthetic ethology approach.

Synthetic Ethology Contrasted with Simulation

The techniques of synthetic ethology must be carefully distinguished from simulation or mathematical modeling. In a simulation, an attempt is made to imitate in a computer or other modeling system the salient aspects of a complex situation that exists, at least potentially, in the real world. The design of a simulation is heavily theory-laden and necessarily highly selective. This is true even for models based on current theoretical and empirical understanding of the phenomena being studied. For out of the multitude of features in the natural situation, only a small fraction can be selected for modeling. This is the Achilles heel of simulation, for an inappropriate selection vitiates the relevance of the model. This problem is especially critical in ethology, because animals respond so sensitively to their environments that it is often unclear whether a feature is relevant or not. Indeed, whether a simulation and its underlying assumptions is considered useful or valid is often based on how robustly it matches our expectations (cf. Burghardt 1984, pp.13-15).

In synthetic ethology, in contrast, we do not attempt to model any existing natural system. Instead we construct “synthetic worlds” which are simple, but complete, and which manifest the phenomena of interest. By “complete” we mean that these worlds define all the conditions that determine whether or not the synthetic organisms “survive,” that is, persist as definite structures. In such a world, then, we can let evolution take its natural course. However, because the world is synthetic, it has important advantages: (1) evolution proceeds much more quickly than in natural worlds; (2) synthetic worlds, though complete, are much simpler than natural worlds; (3) all or most of the variables can be directly controlled; (4) in particular, synthetic ethology permits precisely controlled experiments in which exactly the same initial population and environment can be observed to evolve under different conditions. Thus synthetic ethology permits a degree of experimental control usually seen only in the physical sciences, but without the sacrifice of ecological validity that comes from studying animal behavior in unnatural environments.

The direct inspiration for our work is the *synthetic psychology* of Valentino Braitenberg (1984). He observes that the “law of uphill analysis and downhill invention” (p. 20) means that it will generally be much more difficult to analyse an existing complex mechanism to determine its internal structure than to design a system of comparable behavior. The reason, in essence, is that deduction is usually easier than induction. The lesson is that traditional analytic (empirical) approaches to the study of behavior should be complemented by synthetic approaches, which, if nothing else, may suggest hypotheses to guide analysis and observation. By affording more direct experimental control, the synthetic approach seems more likely to suggest behavioral laws of great generality, which can then be tested by the traditional approaches.

Synthetic Ethology applied to Communication

The synthetic approach is not only appropriate for strictly behavioral approaches to phenomena, but is especially appropriate for phenomena, such as meaningful communication and language, that philosophers take as examples of *intentionality* (e.g., Dennett 1987; Bekoff & Allen 1992).

Although some might hold that a scientific understanding of such mental phenomena requires an

account in terms of the underlying neurophysiology (MacLennan 1988), the evolution of the complex interaction of nervous systems and their environments is well beyond the current capabilities of science and even basic terminology are still in dispute (e.g., Guilford & Dawkins 1991; Blumberg & Alberts 1992). Further, in order to preserve the pragmatics of communication, it is necessary to investigate it in the organisms' *natural* environment, that is, the environment to which it had become coupled through natural selection (MacLennan 1990, 1992). To achieve this coupling while retaining experimental control, we decided to have synthetic organisms evolve in simple, yet complete synthetic worlds. Thus the origin of the term for our approach, *synthetic ethology*.

The advantages of synthetic ethology can be illustrated by the difficulties inherent in the behavior of concern here: communication. The first problem is one of definition: how can a behavior be identified as a communication act, an issue of some controversy (e.g., Burghardt 1970; Slater 1983)? It is not sufficient to say it is a behavior of one organism that influences the behavior of another, since such a definition is much too broad; it makes almost all behaviors communication. On the other hand, if we state that the communicator must have some intention of influencing the receiver's behavior, then our definition depends on the poorly understood and teleological notion of intent. Indeed, it begs the question of intentionality, which it was our purpose to investigate.

A definition by Burghardt (1970) finessed the issue of intent by the requirement that the behavior be likely to influence the receiver in a way that benefits, in a probabilistic manner, the signaller or some group of which it is a member. Often this group is based on genetic relatedness, and thus the communication act must be adaptive, in an evolutionary sense, to the communicator. One difficulty with this definition is that it may be difficult to establish operationally whether in fact a particular behavior, on the average, enhances the fitness of the signaller or its group, since the benefit could be long term or indirect, and confounded with many other influences. Here is a situation where synthetic ethology may be helpful, for we can construct two synthetic "worlds" differing only in whether the behavior in question is possible. Then we can follow the course of

evolution in these worlds and observe any differences that arise. Furthermore, we can run the experiment many times, with controlled or random variations, and so separate general laws from historical accidents. In ethology, the most comparable approach would be the use of “natural experiments,” in which functions can be inferred from comparative studies of animals with different ecologies (cf. Tinbergen 1963).

Related Methodologies

Synthetic ethology can be compared with several related techniques. One is *computational neuroethology*, which is concerned with simulating the neural mechanisms underlying an organism’s interaction with its environment (Beer 1990; Beer et al. 1990; Cliff 1990). To date, most research in computational neuroethology has focused on single organisms in simple environments, but when it investigates populations in complete environments then it corresponds to synthetic ethology.

Artificial life is a very new discipline (Langton 1989; Meyer & Wilson 1991; Langton et al. 1992); it has been defined as “a field of study devoted to understanding life by attempting to abstract the fundamental dynamical principles underlying biological phenomena, and recreating these dynamics in other physical media — such as computers — making them accessible to new kinds of experimental manipulation and testing” (Langton 1992, p. xiv). In addition to behavioral phenomena, artificial life is concerned with prebiotic chemical evolution, self-reproduction, artificial metabolism, evolutionary dynamics, development, learning and cultural evolution.

Synthetic ethology holds much the same relation to AL (artificial life) as synthetic psychology holds to AI (artificial intelligence). Both AL and AI are quite broad in their scope, but emphasize the construction of systems exhibiting the relevant phenomena (life or intelligence). In both synthetic ethology and synthetic psychology, in contrast, the concern is more specifically scientific rather than technological, since the goal is to use the synthetic approach to understand the natural phenomena, rather than to produce useful artifacts. Nevertheless, there is a great deal of overlap,

and when AL is used to study behavioral and social phenomena that are closely coupled to their environment, then it is essentially the same as synthetic ethology.

Finally, we need to contrast our analysis of communication with studies of the evolution of cooperation (e.g., Axelrod 1981, 1984). As described below, our experiments investigated the evolution of communication in an environment that selects for cooperation. As in the prisoner's dilemma games, there is a cost to attempting to cooperate; typically it is a lost-opportunity cost, but we have also imposed specific communication costs, with little difference in outcome. (By a lost-opportunity cost we refer to the fact that, although the reward for cooperation may be high, there is no reward for an attempted cooperation that is not reciprocated, whereas noncooperative action has a reliable, though probabilistic, reward.) Thus, while the evolution-of-cooperation experiments differ from ours in that they have not addressed the evolution of communication, they are similar in that they have shown that cooperative behavior can evolve if it leads to higher fitness than noncooperative behavior.

Experiments

General Methods

Prerequisites to Communication

Our goal in these experiments was to design a synthetic “world” that is as simple as possible while still permitting communication to evolve. A first prerequisite for communication is that some organisms have access to information (knowledge) that others do not, for if they all have access to the same information, no communication is necessary. The nonshared information could be about the organism's own internal state (e.g., hunger), or it could be about features of the external state of the environment that cannot be directly perceived, or as well perceived, by the other organisms.

A second requirement is that this local information must “matter”; it must have some relevance or significance to the organisms. In particular, if its transmittal is to constitute communication according to the definition given above, then its reception must tend to confer some advantage on

the sender or its group (which may, of course, include the receiver). An example of such communication is that which facilitates cooperation that benefits both the sender and the receiver. Therefore we have constructed our synthetic world to select for a kind of cooperation that can be accomplished more effectively by access to information that is not directly perceivable, but which might be communicated.

Shared and Local Environments

For communication to be possible there must be a shared *global environment* in which some organisms can make changes that can be sensed by other organisms. We made this shared environment as simple as possible: a single variable that can take on one of a finite number of values. It is easiest to think of this variable as a medium such as the “air” and each possible value as one of a finite number of distinct “sounds,” only one of which is allowed to be in the air at a time. Notice that the shared environment has no geometry: there is no space through which signals propagate, nor do they have a direction; when a signal is emitted it is immediately available to every simulated organism.

For communication to have any value, some organisms must have access to information that others could use, but which is otherwise unavailable to the second organism. Therefore, in these experiments, we gave each simulated organism a *local environment* which it alone was able to sense directly. The local environment can be thought of as some features of the immediate vicinity of the organism, which can be “felt” by that organism, but not “seen” by any more distant one. Since organisms cannot sense each other’s local environments, there are no geometrical relations between them. They are not in a rectangular grid, nor are some closer than others. The implications of this local “privacy” and the lack of geometrical assumptions can be the subject of future studies.

In these experiments we made the local environments as simple as possible: each was a single variable that could take on one of a finite number of values. We often find it useful to think of a local environment’s state as representing some kind of potential prey in the organism’s vicinity.

Cooperation

In these experiments, cooperation takes the form of one organism acting in a way appropriate to a different organism's local environment. We may think of it as two animals together capturing a large prey animal that neither could bring down alone. Since local environments are set randomly and a potential cooperator cannot directly sense the state of another's local environment, it has only two choices: to guess the other's local state, or to acquire it by means of communication (through the global environment). We can therefore detect communication by levels of cooperation that exceed that expected by chance.

In most of our experiments we have found it useful to impose an additional requirement on cooperation. With a limited number of possible local environment states (typically eight to sixteen) and moderate size populations (typically at least 100 organisms), it is virtually assured that any action attempted by a potential cooperator will match the local environment of some other organism. Therefore, we require a potential cooperator's action to match the local environment of the last emitter. This places more selective pressure against guessing (and thus in favor of communication) and therefore speeds up evolution; it does not have other significant effects.

The point is worth repeating that in synthetic ethology we are not attempting to simulate any specific natural system. Rather, since we are creating a world from scratch, we may build into it any "natural" laws convenient for studying the phenomena of interest. Thus, while there is pedagogic value in conceptualizing the cooperative rule in terms of "two animals together capturing a large prey animal," the fact is that we have constructed the world so that the organisms in it will be likely to reproduce only if they "cooperate" in the specified way.

Naturally, if the synthetic world is too alien, we may doubt the applicability to our world of any observations made of the synthetic world. This, however, is an unavoidable pitfall of experimental science. In constructing an artificial situation to facilitate the experiment, we run the risk of altering precisely the phenomena of interest. Replication and varying conditions will, however, allow assessment of such a possibility.

Simulated Organisms

We call the simulated organisms that were used in these experiments *simorgs*. There are a number of ways to control the behavior of simulated organisms. An obvious and biologically plausible method would be by a simulated neural network, and we have used that approach (MacLennan & al. 1990), as have other researchers (Werner & Dyer 1992). In these experiments, however, we chose a different control mechanism, called a *finite-state machine* because it can be in only one of a finite number of states. In effect it has only a finite number of bits of memory. Finite-state machines have the advantages that they are both readily understood intuitively and easy to represent in genetic strings for simulated evolution.

The behavior of a finite-state machine is determined by a number of condition/effect rules. The conditions include the *internal state* of the simorg (corresponding to its “short-term memory,” “mental state,” or “physiology”) as well as the sensible state of its environment (for us, the states of the global environment and of its own local environment). The effects include the new internal state of the simorg, as well as any responses it makes, such as emitting a signal or attempting to cooperate with another simorg. In these experiments the rules have the form:

$$(\Sigma, \gamma, \lambda) \Rightarrow (\Sigma', R),$$

where Σ is an internal state value, γ is a global state value, λ is a local state variable, Σ' is a new internal state value, and R is a response. We may paraphrase the rule: “If my current internal state is Σ , and I sense γ in the global environment, and the situation in my local environment is λ , then my new internal state will be Σ' and I will make response R .” In these experiments the response is either an *emission* (*signal*) or an *action*. A emission has the form:

$$\text{emit}(\gamma'),$$

and puts the global environment into state γ' . An *action* has the form:

$$\text{act}(\lambda'),$$

and represents an attempt to cooperate with a simorg in situation λ' . If the local environment of the last emitter is in state λ' then they cooperate, otherwise they do not. It must be stressed that $\text{act}(\lambda')$ does not *do* anything beyond causing λ' to be compared to the last emitter's local environment, which results in cooperation if it matches, and a consequent increase in fitness (described later).

This rule for cooperation is not intended to model any specific situation in the natural world. Our hypothesis relates communication and cooperation, but it says nothing about the kind of cooperation. In synthetic ethology, since we construct a "world" for the purposes of the experiment, we have the freedom to build in a "law of nature" that selects for any sort of cooperation we choose. Nevertheless, the following analogy may help the rule to be understood. We can imagine a group of hunters stationed in different territories. When prey enters a hunter's territory he or she may signal that fact, and another hunter may respond and help to bring down the prey, but to cooperate successfully, the helper must bring the appropriate tool (e.g., a fishing line, a bola, or a bow and arrow).

Finite-state machines have a rule for every possible condition: in these experiments, for every possible combination of internal state, global state, and local state. Thus the simorg's behavior is completely determined. The collection comprising all a finite-state machine's behavioral rules is called its *transition table*, because it defines the machine's transition from state to state.* The transition table can be depicted by a diagram such as Fig. 1, which shows its internal states (the circles) and the conditions on some of the transitions between them (the lines). The machines used in these experiments had from 64 to 128 rules in their transition tables, so they are too complicated to show in this way.

* A transition *table*, which is a one-dimensional array that defines the behavior of a deterministic machine in each possible situation, must be distinguished from a transition *matrix*, which is a two-dimensional array that defines the probability of state transitions for a nondeterministic machine.

A transition table can be represented in a “genetic string” in a simple way. We have seen that the global environment states can be represented by integers $1, 2, \dots, G$ and the local environment states by $1, 2, \dots, L$. If the machines have I internal states, then these too can be represented by integers $1, 2, \dots, I$. Therefore, we see that a transition table must have IGL entries to define the simorg’s behavior under all conditions. Thus a simorg’s complete behavioral repertoire can be defined by a genetic string with IGL “genes,” one for each possible condition. For example, in the first series of experiments described below, $L = G = 8$ and $I = 1$, and so each simorg had 64 genes. In the second series of experiments, $L = 8$ and $G = I = 4$, so these machines had 128 genes.

The effect of each rule is to determine one of I possible new internal states, and to either emit one of G possible signals, or to attempt one of L possible actions. Thus there are $E = I(G + L)$ possible effects for each condition. These can be encoded in the genetic string if each gene has E alleles, represented by the integers $1, 2, \dots, E$. In these experiments E was typically 16 or 64 (i.e., 4 to 6 bits). The use of genetic strings is described below (“Selection and Reproduction”).

Finally, we turn to learning and experiential effects, broadly considered. These are an important factor in the behavior of many organisms, and one of our goals is to use synthetic ethology to study it. Learning is especially relevant to communication, since we would like to compare the efficacy and flexibility of innate (or closed) and learned (open) communication protocols (“languages”). Guilford & Dawkins (1991) point out that learning and memory are key, but little studied, processes in signals and their evolution. In these experiments we studied only the simplest kind of single-case learning. Specifically, when a learning simorg “makes a mistake” by acting noncooperatively, then we change its transition table so that it *would have* acted correctly under the current conditions. That is, suppose the rule that matches the current state is:

$$(\Sigma, \gamma, \lambda) \Rightarrow (\Sigma', \text{act}(\lambda')),$$

but that the local environment of the last emitter is in state $\lambda'' \neq \lambda'$, and so cooperation fails. The acting simorg learns by changing this rule in its transition table to:

$$(\Sigma, \gamma, \lambda) \Rightarrow (\Sigma', \text{act}(\lambda')),$$

which would have been the correct action. The assumption is that there is some chance that it will be the correct action the next time these conditions occur. Clearly, this learning rule could easily lead to instability. Nevertheless, it seems to work rather well (see below, “Results”).* It may be conceived as similar to an animal using observational information to correct its response after the fact from a limited set of alternatives. Alternately, we may imagine a hunter signalling that a certain kind of prey is in his or her territory; potential cooperators must bring the appropriate tool in order to cooperate, but if they do not, then they may alter their behavior once they have discovered what tool they should have brought.

Selection and Reproduction

Whenever a simorg acts in a way appropriate to the local environment of the last emitter, then we say that cooperation has taken place, and award both the emitter and the receiver a point of credit; we may think of it as the stored energy resulting from eating food. We call the credit accumulated by a simorg its “fitness,” but it must be kept in mind that this means nothing more than the number of times it has successfully cooperated. Since this “fitness” is used to determine which simorgs “breed” and which “die,” we are selecting for simorgs that cooperate.

An experiment runs for a number (typically several thousand) of *breeding cycles*, each of which comprises a number of *environmental cycles*, each in turn comprising several *action cycles*. During an action cycle each simorg reacts to its environment as determined by its transition table (its behavioral rules). The simorgs react one at a time in a fixed order determined by their position in a table. After five action cycles, all the local environments are changed randomly, and the simorgs are allowed to react again. After 10 such environmental cycles, two simorgs (parents) are

* Instability could result because a single mistake causes a change of behavior. Thus the behavior of a well-cooperating simorg could be disrupted by a signaller that emits the wrong signal. Nevertheless, single-case learning is not uncommon in nature.

selected for breeding and one simorg (moriturum) “dies,” with all selections being based on the “fitness” accumulated during the breeding cycle. The resulting single offspring replaces the “deceased” simorg in the population, thus keeping the population size constant. In effect we are simulating a population equal to the carrying capacity of the environment. This simplifies the experiment, since we don’t have to worry about populations dying out or growing exponentially. At the end of a breeding cycle every simorg’s fitness is reset to zero, and a new breeding cycle begins. The experiment continues for a fixed number of breeding cycles (5000 - 200,000 in these experiments).

The parents and moriturum are selected stochastically, depending on their fitness. We have made the probability of a simorg’s being selected as a parent proportional to its fitness, and its probability of being the moriturum a simple monotonically decreasing function of its fitness (see MacLennan 1990, 1992 for formulas). In effect simorgs with more “energy reserves” are more likely to breed, and those with less are more likely to die. However, since the selection process is stochastic, a very unfit simorg will occasionally be selected as a parent, or a very fit one as the moriturum. (Indeed, the same simorg could be selected for all three — effectively breeding with itself and then dying in childbirth!) However, in most cases the parents will not die, since they probably have high fitness, whereas the moriturum probably has low fitness.

A brief walk through an entire breeding cycle may be helpful. First the local environments are initialized randomly. Next all the simorgs are allowed to react, one at a time and in order. A simorg’s transition table tells it how to react to its environment, comprising the global environment and its local environment. It may react by emitting a symbol, that is, by changing the global environment, or it may react by attempting to cooperate, an action that is successful only if it acts appropriately to the last emitter’s local environment. Thus communication takes the following form. Based on its local environment (and the global environment) a simorg places a symbol in the global environment. At a later time (and before another simorg has emitted a symbol), a second simorg may react to the symbol in the global environment (and to its own local environment) by attempting to cooperate. If its action matches the local environment of the last emitter, then they

both are given a point of fitness, otherwise they are not. In the latter case, if learning is enabled, then the receiving simorg changes its action under those circumstances (local and global environment) to what *would have been* correct (though there is no guarantee that it will be correct under the same circumstances next time).

Five times all the simorgs are allowed to react in this way, then the local environments are reset randomly, and the simorgs are allowed to react to this new set of situations. After ten such changes of the local environments, breeding takes place. Parents are chosen randomly, but preference is given to the simorgs that have cooperated the most since the last breeding. Likewise simorgs that have not cooperated successfully are most likely to be chosen to “die,” that is, to be replaced by the offspring of the parents. After, say, 5000 breeding cycles of this kind, the experiment is terminated and statistics are calculated.

Genotypes and Phenotypes

Associated with each simorg are two packages of information, which we call the *genotype* and the *phenotype*. The phenotype is the transition table used to generate the simorg’s behavior, and the structure that is modified if and when the simorg learns. In contrast, the genotype never changes during the “life” of the simorg, but is related to the genotypes of its parents and offspring by certain *genetic operators* described next.

When the parents have been chosen, the genetic strings representing their genotypes are combined by a process called *two-point crossover*. This means that (1) two points, η and θ , are chosen randomly, and (2) between locations η and θ the genes are copied from one parent, and (3) the remainder of the genes are copied from the other parent (see Fig. 2). There is no distinction between dominant and recessive genes, and all simorgs belong to one sex although two are required for reproduction (hermaphrodites).

With low probability (usually 0.01) the genetic string is mutated after crossover. If a mutation is to occur, then a randomly chosen gene is replaced by a random allele (single point mutation). In effect, a mutation replaces one entire entry in the finite-state machine’s transition table. After the

offspring's genotype has been determined by crossover and mutation, it used to produce the actual phenotypic transition table, and the resulting simorg replaces the moriturum.

Experiments, Measurements, and Statistics

If true communication (involving a sender) rather than information extraction is taking place, then it must “confer some advantage (or the statistical probability of it) to the signaler or his group” (Burghardt 1970, p. 16). This is the kind of condition that synthetic ethology is ideally suited to test, for we can start two synthetic worlds with the same population of random simorgs, but prevent communication in one and not interfere with it in the other. Systematic differences in the future evolution of these worlds can then be attributed to communication. In particular, we *suppress communication* by randomizing the state of the global environment at every opportunity (specifically, after every simorg responds). In effect this raises the “noise level” to the point that no communication is possible; under these circumstances cooperation can result only from lucky guesses.

A second comparison that we can carry out with synthetic ethology is the evolution of communication with and without learning. It is a simple matter to allow the learning mechanism to be turned on or off, a degree of control that would be difficult to achieve with natural organisms. (As explained above, we considered only the simplest kind of single-case learning.)

We gathered a number of statistics during and after the experiments. The most basic was the “fitness,” that is, the number of successful cooperative acts. This count was maintained for each simorg and entered into the breeding and replacement decisions as already described. We also computed the average fitness of the population, smoothed over 50 breeding cycles, which was written to a file for later plotting (e.g., Figs. 3-5). After each experimental run was complete, these data were used for other purposes, such as determining by linear regression the average rate of fitness increase. These measurements provide the most basic way of assessing the functional effect of communication, since we can compare the rates of fitness change and fitness levels

achieved under varying conditions, such as with communication suppressed or not, and learning enabled or not.

Beyond merely detecting the presence of communication, we are also interested in studying its *structure*. In these experiments this was accomplished by gathering statistics on the local and global environment states involved in successful cooperations. Specifically, whenever two simorgs cooperated we incremented a table entry corresponding to the “symbol” γ in the global environment and the “situation” λ in the local environment of the last emitter. The reason is that a successful cooperation is prima facie evidence that for the emitter and receiver the symbol γ denotes the situation λ . Of course, the cooperation could have resulted from a “lucky guess,” but these will be uniformly distributed across the table; frequently occurring symbol/situation pairs indicate structured symbol use. We refer to this table as a *denotation matrix* because it refers to apparent use of symbols to denote situations.

We found that when the denotation matrix was accumulated over the entire experiment the initial, random use of symbols obscured their later more structured use. For this reason we changed the denotation matrix to reflect only the most recent 50 breeding cycles of the experiment. Thus we could track the evolution of the “language” used by the simorgs. The matrices displayed in Tables 1-3 show the statistical structure of the communication at the ends of the experiments (the last 50 breeding cycles).

Although, as will be seen below, the increased structure of the denotation matrix when communication is not suppressed is apparent to the eye, we also quantified its degree of structure. If no real communication were taking place, then symbol/situation combinations should occur with approximately equal frequency; therefore the denotation matrix would be quite uniform. This suggested several ways of measuring the “structure,” or deviation from uniformity, of the matrix. One obvious possibility was the standard deviation, since it measures the spread of a distribution. Unfortunately, the standard deviation varies with the number of cooperations, which makes it

difficult to compare across experiments. For this reason we computed the coefficient of variation (CV) of the matrix, which measures the standard deviation in units of the mean, $CV = \sigma/\mu$.

Another well-known measure of the uniformity of a probability distribution p_k is its *entropy* or *information content*, which is defined:

$$H = -\sum_k p_k \log p_k .$$

The entropy is maximized by a uniform distribution, and decreases as it becomes less uniform. The denotation matrix gives the frequencies of symbol/situation combinations; by treating this frequency table as a probability distribution we were able use entropy to measure its uniformity. It is easy to show that for the uniform distribution the entropy is $H_{\max} = 2 \log N$, where $N = G = L$ is the number symbols or situations (equal in these experiments). (For simplicity of interpretation we use base 2 logarithms throughout.) For comparison, the entropy of an “ideal” language, in which there is a one-to-one correspondence of symbols and situations, has an entropy $H_{\text{ideal}} = \log N$. In these experiments, since $N = 8$, $H_{\max} = 6$ and $H_{\text{ideal}} = 3$.*

There are, of course, other ways besides coefficient of variation and entropy for measuring the structure of the evolved “language.” For example, we have also used chi-square to test the predictability of situations from symbols and vice versa. The conclusions to be drawn are the same, however.

Experiment 1: Evolution of Simple Signals

Methods

In these experiments the number of possible global environment states G was equal to the number of possible local environment states L . That is, there were just enough possible “sounds” to describe the possible “situations.” In particular, $G = L = 8$, so the probability of a correct guess is

* An “ideal” denotation matrix has exactly one nonzero entry in each row and in each column, and all these nonzero entries are the same, $1/N$. Thus the entropy is $-\log(1/N) = \log N = 3$, for $N = 8$.

1/8. The machines were given 0 bits of internal memory; that is, they had only one possible internal memory state. As a result, the rules in the transition tables of the finite-state machines have the form of simple stimulus/response rules, where the stimulus is a symbol/situation pair (γ, λ) , corresponding to states of the global and local environments, and the responses are emissions or actions. An emission rule has the form:

$$(\gamma, \lambda) \Rightarrow \text{emit}(\gamma'),$$

and an action rule has the form:

$$(\gamma, \lambda) \Rightarrow \text{act}(\lambda').$$

Experiments were generally run for 5000 breeding cycles, although some were run 10 times as long. The local environments were changed randomly after every five action cycles, and a simorg was replaced by breeding after every 10 local environment changes. The population size was 100 and the mutation rate was 0.01, meaning that with 0.01 probability a gene was replaced by a random allele. The program itself is about 650 lines of Common LISP code, and runs identically on a variety of workstations (a copy will be provided upon request). Each 5000-cycle experiment took about 100 min. of computer time (270 min. when learning was enabled) on a Sun-4 workstation.

Results

All in all, we made over 100 experimental runs, differing in random seed and other parameters (e.g., suppression of communication). We consider three evolutions starting from the same random initial population; the results described are typical of all the experiments. Figure 3 shows the fitness of a population of nonlearning simorgs when communication was suppressed. (All the fitness plots included in this report are smoothed by a moving average of 50 breeding cycles. This makes the trends much more apparent.) The fitness stayed within 12% of 6.25 cooperations per

breeding cycle, the level of cooperation predicted for a population of “guessing” simorgs.* Linear regression detects a slight upward slope (3.67×10^{-5} coops./br. cyc./br. cyc.), a phenomenon we address later. Consider next Fig. 4, which shows the degree of cooperation among nonlearning simorgs when communication was not suppressed; by the end of 5000 breeding cycles it had reached 10.28 cooperations per breeding cycle, which is 65% above the chance level. Linear regression shows that the fitness is increasing at a rate of 9.72×10^{-4} , which is 26 times as fast as when communication was suppressed. Finally consider Fig. 5, which shows the degree of cooperation when communication was not suppressed and the machines were capable of learning. In this case the fitness achieved was 59.84 cooperations per breeding cycle, which is 857% above the chance level, and was increasing at a rate of 3.71×10^{-3} , which is 3.82 times the rate without learning and one hundred times the rate when communication was suppressed.

Now we turn to the denotation matrices resulting from three evolutions beginning with the same random initial population. When communication was suppressed, the matrix shown in Table 1 resulted. Visually, it is quite uniform, and this is confirmed by the coefficient of variation $CV = 0.68$ and by the entropy $H = 5.66$, which is near its maximum $H = 6$. When communication was not suppressed, a much more structured denotation matrix resulted (Table 2), an observation confirmed by a higher coefficient of variation, $CV = 2.04$, and a lower entropy, $H = 3.95$; Fig. 6 shows how the entropy decreased during the evolution of the language. Finally, when learning is also enabled, we obtained the denotation matrix in Table 3, which is slightly more structured ($CV = 2.72$, $H = 3.47$).

By looking at the denotation matrices, which represent communicative activity over the last 50 breeding cycles, we can draw some conclusions about the “language” used by the simorgs. First, some symbols have almost come to denote a unique situation; for example Table 2 shows that 90% of the recent uses of symbol 5 (in successful communication acts) referred to situation 2. Second, certain situations have symbols that typically denote them; in over 99% of the recent

* Analysis and calculation of the expected fitness of “guessing” simorgs can be found in MacLennan (1990).

communications involving situation 7, it was denoted by symbol 2. On the other hand, the matrix also shows ambiguous symbol use; for example, in 65% of its recent uses symbol 2 denoted situation 5, and in 31% it denoted situation 7. Since the matrix reflects the collective communication acts of the entire population, we cannot tell from it whether this multiple use of symbols results from two subpopulations, or from individual simorgs using the symbol to denote two situations. Of course this information can be determined by looking at the individual simorgs, a topic to which we now turn.

One of the significant advantages of synthetic ethology, beyond the experimental control afforded, is the ability to have complete access to the mechanism by which behaviors are generated. There need be no mystery about how the simorgs are communicating, because the process is completely transparent. For example, we have “dissected” the most fit simorg from the final population of an experiment to see how its use of symbols compares with that of the population; we found over twice the number of coincidences that would be expected by chance (MacLennan 1990, sect. 3.3). Since the simorgs’ structures are available in the computer, it is possible to write programs that scan them to look for patterns or to compile statistics.

To illustrate some of the interesting results of synthetic ethology experiments, Fig. 7 shows the average fitness for an experiment comparable to Fig. 5 (i.e., communication permitted, learning enabled), but with a different random starting population. The fitness increased for the first 3000 breeding cycles at a rate of 9.77×10^{-4} , which is a little slow; it’s more typical of a nonlearning population, such as shown in Fig. 4, than of a learning population. At approximately $t = 3000$ the fitness decreased rapidly. One possible cause is a “genetic catastrophe,” that is, a chance coincidence of unfortunate occurrences, such as highly fit simorgs dying and unfit simorgs breeding. In other words, the average fitness could decrease because, coincidentally, the fittest simorgs died and the least fit simorgs survived. Another possible cause is *language instability* resulting from the very simple learning rule that we use; a chance encounter between an effective communicator and noncommunicators can cause the communicators to rapidly “forget” their

language.* In other words, the fittest simorgs could become much less fit by losing their ability to cooperate effectively. In any case, about time $t = 4000$ the population began to recover, and at a quite rapid rate, 5.08×10^{-3} , which is 37% faster than that observed in Fig. 5.

Experiment 2: Evolution of Complex Signals

Methods

In these experiments the number of possible global environment states was less than the number of local environment states, $G < L$. Thus, an adequate description of a local environment situation would require the use of two or more symbols. We picked $L = 8$, $G = 4$, so that two symbols would be more than adequate (since with $G = 4$, two symbols could describe up to 16 situations). So that the simorgs would be able to remember which of the four symbols they had already sensed, we gave them two bits of memory. Thus, the number of possible internal or “mental” states was $I = 4$. Notice that the simorgs still produce only one symbol per emission, and therefore at most one symbol per action cycle. However, since they have the theoretical ability to remember the last symbol they emitted, they have the theoretical capability of emitting coordinated pairs of symbols (on successive action cycles). Similarly, since they have the theoretical ability to remember the previous symbol they received, they likewise have the theoretical capability of recognizing strings of symbols of length two. However, remembering previous symbols is not “wired into” our simorgs, though it may evolve.

As in the previous series of experiments, a breeding cycle comprised 10 sets of local environment states, which were changed randomly every five action cycles. As before, the population size was 100 and the mutation rate was 0.01. Due to the increased complexity of the

* Of course, the cause of the fitness decline need not be a mystery; synthetic ethology permits us to “replay” that period at “slow speed,” that is, to monitor the individual events as they occur. We have not investigated this particular instance of fitness decline since it was an isolated occurrence and did not seem relevant to our study.

machines, evolution was allowed to continue much longer, from 1×10^4 to 2×10^5 breeding cycles.

In order to increase selective pressure we also tried several variations, including raising the fitness to a power (to exaggerate the difference between cooperators and noncooperators) and imposing a penalty for unsuccessful cooperation (in which case we also provided initial “free” fitness points — analogous to an initial energy reserve). We also tried the experiment with learning simorgs.

The program in this case was coded in both Common LISP (740 lines) and FORTRAN 77 (785 lines), and runs on a variety of workstations. The 10000-cycle experiments described here used the FORTRAN version (available upon request), and ran about 20 min. on a Sun-4 workstation.

Results

A typical experiment ran for 10000 breeding cycles. The population size was 100, and the machines had $I = 4$ internal states. The local environments had $L = 8$ states, but the global environment had only $G = 4$ states. The fitness was squared and an initial fitness of 10 points (equivalent to 10 cooperations) was provided. The penalty for unsuccessful cooperation was the same as the reward for successful cooperating, one point (thus each mistake cancels a cooperation). Thus the differences from experiment 1 were: (1) there were 4 internal states rather than 1 (thus permitting memory); (2) there were 4 global environment states rather than 8 (thus there were insufficient states to denote the 8 local environment states); (3) the fitness computation was slightly different.

Table 4 shows a typical denotation matrix. By the end of the experiment, the entropy had decreased from its maximum of $H_{\max} = 7$ to $H = 4.62$, which may be compared with the entropy of the “ideal” matrix, $H_{\text{ideal}} = 3$. (See also Fig. 8.)

This denotation matrix reveals several characteristics of the language that had evolved. First, nonzero entries tend to occur in blocks of four, which reflects the fact that the second (most recent)

symbol is usually the significant one. For example, situation 7 is denoted exclusively by “sentences” of the form ‘X2’. Similarly, situation 5 is denoted primarily by ‘X1’ sentences. This dependence on the second symbol is not surprising, since it doesn’t require the simorg to remember the previous symbol. We do, however, observe some situations in which the both symbols are being used. For example, while ‘00’ and ‘01’ often mean situation 4, the sentence ‘20’ mostly denotes situation 0 and only infrequently denotes situation 4. Other experiments showed a preponderance of sentences of the form ‘XX’, which is another way of avoiding the need for memory.

We conclude that in these experiments the machines did not evolve to make full use of the communicative resources to denote the eight possible situations. That this was not a simple matter of insufficient time is shown by Fig. 8, which shows that the entropy had stopped decreasing after about 5000 breeding cycles. The population reached an evolutionary plateau. This indicates that making the step to multiple-symbol syntax is evolutionarily hard, an hypothesis that will be explored in future experiments.

General Discussion

We summarize the main results of the two series of experiments. In the first we showed that under a variety of conditions a population of simple machines can evolve the ability to communicate in an environment that selects for cooperative behavior. When the possibility of communication is suppressed, the machines cooperate at a very low level, as predicted by theoretical analysis. On the other hand, when communication was not suppressed, the machines did communicate, and as a result they achieved a much higher level of cooperation, and the rate of fitness increase was an order of magnitude greater. Giving the machines a very simple kind of single-case learning allowed them to cooperate even more effectively and to increase their fitness at a significantly greater rate. We also investigated the regularity of apparent symbol use and found it was nearly random when communication was suppressed, but quite organized (as measured by entropy) when it was not suppressed. While a thorough comparison of our results to biological phenomena is not

possible here, we do note that simple communication appears to occur in protista, bacteria, and among cells in multicellular organisms, and may have been critical in the early development of living systems. The marked advantage of learning in communication and symbolic communication is recognized in current cognitive ethology and comparative cognition (e.g. Ristau 1991; Roitblatt 1987; Savage-Rumbaugh, et al. 1993)

These results are quite replicable with different random initial populations, although of course the exact measurements vary. For comparison recall that in the experiments described (Figs. 3-5) the rates of fitness change under the three conditions (communication suppressed, communication permitted, communication with learning) were 0.367, 9.72 and 37.1, respectively, per 10^4 breedings. In two other series of runs with different initial populations the corresponding rates were somewhat lower: 0.16, 8.2, 23.1, and 0.56, 9.4, 29.3. Averaging the three series together shows that the fitness increased 25 times as fast when communication was permitted, and 82 times as fast when learning was also enabled.

The second series of experiments was similar to the first, except that we created a situation in which unique denotation of situations would require the use of pairs of symbols. The results were qualitatively similar, but the population did not unambiguously demonstrate the ability to use pairs of symbols. Additional experiments will be required.

We made several other observations that warrant mention. In all of our experiments in which communication was suppressed and learning was disabled, there was nevertheless a slight upward trend in fitness. We hypothesised that this resulted from the simorgs evolving to take advantage of a specific feature of our selection algorithm. Specifically, we consider cooperation to have taken place if the a simorg's action is appropriate to the situation of the last emitter. Therefore, the simorgs can "beat the odds" by signalling in only a subset of their situations, provided they also evolve to guess only actions in that same subset. We call this phenomenon *partial cooperation*

*through coadaptation.** Several experiments confirmed this hypothesis, but we have not tried to eliminate it since it is a low level effect and its prevention would complicate the experiments (see MacLennan 1992 for details).

Another observation relates to our nondeterministic selection of parents. Recall that the fitness of a simorg determines only the probability of it reproducing; occasionally unfit simorgs will reproduce and fit simorgs will not. An earlier series of experiments used a deterministic strategy: only the two most fit simorgs were allowed to breed. In these experiments we never observed the evolution of communication; the result was always a population of “silent guessers.” We conjecture that in the early stages of the evolution of communication, attempts to communicate are often less effective than guessing, and so the deterministic selection strategy never permitted nascent communicators to reproduce. We have not attempted to test this conjecture, however. But our approach should allow comparison of the various views of the role of natural selection in evolution (cf. Endler 1986) as well as chance environmental events, genetic bottlenecks, and shifting balance scenarios.

Another series of experiments has also demonstrated the evolution of communication in a synthetic world. Werner and Dyer (1992) put selective pressure on communication by making it necessary for effective reproduction. Specifically, the female simorgs are immobile, but can sense the location of the males, whereas the males are blind but mobile. Although males can encounter and mate with females by a random walk, they can do so more effectively if the females “tell” the males how they can be reached. The experiments of Werner and Dyer may be contrasted with ours in several respects. First, they used a neural network model of behavior, whereas we used a finite-state machine model. Second, they used a concrete, “natural” activity (mating) as the source of selective pressure, whereas we took advantage of the flexibility of synthetic ethology and defined a

* In previous papers we used the term “pseudo-cooperation.” Partial cooperation is more accurate, since the cooperation is real, but is possible in only a limited set of situations. Note, however, that no communication is involved, since the symbol placed in the global environment is irrelevant.

more abstract criterion of cooperation. Although their experiments are more biologically suggestive, our design is more amenable to analysis and experimental control. In spite of these differences, many of their observations are qualitatively similar to ours, which lends support to the thesis that synthetic ethology will aid the discovery of general laws.

Neural networks are an alternate to finite-state machines as models of behavior. They have been used successfully in synthetic ethology by Werner and Dyer (1992) and in unpublished experiments of our own (MacLennan & al. 1990). We used finite-state machines in these experiments because of our prior experience in evolving them, and because they can be more easily “dissected,” that is, analysed to determine their structure. This is critically necessary in the early stages of synthetic ethology, although we expect that in the long run neural networks will prove necessary. In addition to their more realistic characterization of the decision making process of vertebrates, neural networks are able to behave more flexibly, especially in the presence of noise. Further, they should permit the investigation of *typical intensity* (Wiley 1983, Morris 1957), the emergence of discrete symbols from continuously variable signals, a critical problem in cognitive science (MacLennan 1993, 1994, in press). Nevertheless, we have shown that even in this simple synthetic world, communication may evolve that exhibits some of the richness of natural communication.

Acknowledgements

This research was supported in part by Faculty Research Awards (1988, 1989) from the University of Tennessee, Knoxville. Members of the Cognitive Science Laboratory and the Ethology program, both at UTK, gave helpful advice during the design of the experiments.

Figure Captions

Fig. 1. Example finite-state machine. At any time the machine is in one of the *states* indicated by the open circles. Environmental/internal condition c_k (which could include sensory input), causes the machine to perform action a_k and move to another state, as indicated by the labeled arcs. Since a machine's current state is determined by its history of prior conditions, it is a kind of memory. In a finite-state machine there are only a finite number of possible states, in effect, a finite number of *bits* of memory, since n bits of memory are equivalent to 2^n states.

Fig. 2. Effect of two-point crossover operation, which is a simplified model of crossover during diplotene stage of meiosis. The genes between η and θ are taken from one parent and the remainder from the other.

Fig. 3. Average fitness of population with communication suppressed and learning disabled. The fitness is observed to vary around the level expected by chance, 6.25. Linear regression shows a slight upward slope, 3.67×10^{-5} .

Fig. 4. Average fitness of population with communication permitted but learning disabled. Linear regression shows that fitness is increasing at a rate of 9.72×10^{-4} , which is 26 times as fast as when communication was suppressed.

Fig. 5. Average fitness of population with communication permitted but learning enabled. Linear regression shows that fitness is increasing at a rate of 3.71×10^{-3} , which is 3.8 times the rate without learning and one hundred times the rate when communication was suppressed. Notice that the average fitness of the population begins at a much higher level than when learning was disabled (Figs. 3, 4); this is because each individual has four more chances within an environmental cycle to respond correctly after it has "learned from its mistake."

Fig. 6. Entropy (measure of disorganization) of the “language” used by the population. (By “language” we mean the communication pattern as represented in denotation matrix.) The maximum possible entropy is 6, which represents completely disordered use of the symbols. Lower entropy is associated with more structured use of symbols. An entropy $H = 3$ is associated with univocal symbol use, that is, each symbol denotes a unique situation. In this case the the language has evolved from a completely disordered state to $H = 3.95$, a more nearly ideal state.

Fig. 7. Average fitness of another population with communication permitted and learning enabled. Notice the sudden decrease of average fitness about $t = 3000$. This might have resulted from a “genetic catastrophe” or from a sudden instability in the language. Prior to that time fitness had been increasing at a rate of 9.77×10^{-4} , which is typical of a population of nonlearners (e.g. Fig. 4). After recovery began at about time $t = 4000$ the fitness increased at a rate of 5.08×10^{-3} , which is 37% faster than the population of learners shown in Fig. 5.

Fig. 8. Entropy of two-symbol “language” used by population. The maximum possible entropy is 7 in this case.

References

- Axlerod, R. (1981). The emergence of cooperation among egoists. *American Political Science Review* **75**: 306-318.
- Axlerod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Beer, R. D. (1990). *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*. San Diego: Academic Press.
- Beer, R. D., Chiel, H. J., and Sterling, L. S. (1990). A biological perspective on autonomous agent design. *Journal of Robotics and Autonomous Systems* **6**: 169-186.
- Bekoff, M., & Allen, C. (1992). Intentional icons: Towards an evolutionary cognitive ethology. *Ethology* **91**: 1-16.
- Blumberg, M. S., & Alberts, J. R. (1992). Functions and effects in animal communication: Reactions to Guilford & Dawkins. *Animal Behaviour* **44**: 382-383.
- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. Cambridge: MIT Press.
- Burghardt, G. M. (1970). Defining 'communication'. In J. W. Johnston Jr., D. G. Moulton and A. Turk (Eds.), *Communication by Chemical Signals*. New York: Century-Crofts.
- Burghardt, G. M. (1984). On the origins of play. In P. K. Smith (Ed.), *Play in Animals and Humans*. Oxford: Basil Blackwell.
- Cliff, D. T. (1990). *Computational Neuroethology: A Provisional Manifesto* (Cognitive Science Research Paper CSRP 162). Sussex: University of Sussex.
- Dennett, D. (1987). *The Intentional Stance*. Cambridge: MIT Press.
- Endler, J. A. (1986). *Natural Selection in the Wild*. Princeton: Princeton University Press.

- Guilford, T., & Dawkins, M. S. (1991). Receiver psychology and the evolution of animal signals. *Animal Behaviour* **42**: 1-14.
- Langton, C. G. (Ed.) (1989). *Artificial Life*. Redwood City: Addison-Wesley.
- Langton, C. G. (1992). Preface. In C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.), *Artificial Life II*. Redwood City: Addison-Wesley.
- Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S. (Eds.) (1992). *Artificial Life II*. Redwood City: Addison-Wesley.
- MacLennan, B. J. (1988). Causes and intentions. *Behavioral and Brain Sciences* **11**: 519-520.
- MacLennan, B. J. (1990). *Evolution of communication in a population of simple machines* (Technical Report CS-90-99). Knoxville: Computer Science Department, University of Tennessee.
- MacLennan, B. J. (1992). Synthetic ethology: An approach to the study of communication. In C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.), *Artificial Life II*. Redwood City: Addison-Wesley.
- MacLennan, B. J. (1993). Characteristics of connectionist knowledge representation. *Information Sciences* **70**: 119-143.
- MacLennan, B. J. (1994). Continuous symbol systems: The logic of connectionism. In D. S. Levine and M. Aparicio IV (Eds.), *Neural Networks for Knowledge Representation and Inference*. Hillsdale: Lawrence Erlbaum.
- MacLennan, B. J. (in press). Image and symbol: Continuous computation and the emergence of the discrete. In V. Honavar and L. Uhr (Eds.), *Artificial Intelligence and Neural Networks: Steps Toward Principled Integration, Volume I: Basic Paradigms; Learning Representational Issues; and Integrated Architectures*. Cambridge: Academic Press.

- MacLennan, B. J., Jerke, N., Stroud, R., and VanHeyningen, M. D. (1990). *Neural Network Models of Cognitive Processes – 1990 Progress Report* (Technical Report CS-90-125). Knoxville: Computer Science Department, University of Tennessee.
- Meyer, J.-A., and Wilson, S. W. (1991). *From Animals to Animats*. Cambridge: MIT Press.
- Morris, D. (1957). 'Typical intensity' and its relation to the problem of ritualization. *Behaviour* **11**: 1-12.
- Ristau, C. A. (Ed.) (1991). *Cognitive Ethology: The Minds of Other Animals*. Hillsdale: Lawrence Erlbaum.
- Roitblatt, H. L. (1987). *Introduction to Comparative Cognition*. New York: Freeman.
- Savage-Rumbaugh, E. S., Murphy, J., Sevcik, R. A., Brakke, K. E., Williams, S. L., and Rumbaugh, D. M. (1993). *Language Comprehension in Ape and Child*. Monographs of the Society for Research in Child Development **58**: 3-4, Serial No. 233.
- Slater, P. J. B. (1983). The study of communication. In T. R. Halliday and P. J. B. Slater (Eds.), *Animal Behavior Volume 2: Communication*. New York: W. H. Freeman.
- Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für Tierpsychologie* **20**: 410-433.
- Werner, G. M., and Dyer, M. G. (1992). Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.), *Artificial Life II*. Redwood City: Addison-Wesley.
- Wiley, R. H. (1983). The evolution of communication: Information and manipulation. In T. R. Halliday and P. J. B. Slater (Eds.), *Animal Behavior Volume 2: Communication*. New York: W. H. Freeman.

Tables

Table 1. Denotation matrix resulting from experiment with communication suppressed and learning disabled*

situation →		0	1	2	3	4	5	6	7
symbol ↓	0	180	201	27	712	149	296	254	292
	1	202	191	21	707	140	268	240	338
	2	196	199	24	699	145	284	235	290
	3	168	154	20	713	135	312	214	314
	4	200	182	15	643	149	310	226	284
	5	206	183	28	684	142	283	243	280
	6	204	191	21	676	145	290	221	310
	7	198	186	19	689	128	276	236	297

* A denotation matrix shows the symbol/situation combinations occurring in successful communications during the last 50 breeding cycles of the experiment. It is thus a “dictionary” for the “language” used by the simulated organisms.

Table 2. Denotation matrix resulting from experiment with communication permitted and learning disabled.

sym. sit.	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

Table 3. Denotation matrix resulting from experiment with communication permitted and learning enabled.

sym. sit.	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. Denotation matrix resulting from experiment investigating the evolution of complex symbols*

sit. 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

* “Symbol” denotes the last two symbols emitted by a simulated organism involved in a successful communication act.

Figures

<figure *FSM*>

<figure *Crossover*>

<figure *Fitness (920821)*>

<figure *Fitness-c (920821a)*>

<figure *Entropy-c (920821a)*>

<figure *Fitness-cl (920822c)*>

<figure *Fitness-cl (920629c, d)*>

<figure *Entropy (901115e = 9204106)*>