

Synthetic Ethology: An Approach to the Study of Communication*

Bruce MacLennan
Computer Science Department
University of Tennessee
Knoxville, TN 37996-1301
internet address: maclellan@cs.utk.edu

November 23, 1990

Abstract

A complete understanding of communication, language, intentionality and related mental phenomena will require a theory integrating mechanistic explanations with ethological phenomena. For the foreseeable future, the complexities of natural life in its natural environment will preclude such an understanding. An approach more conducive to carefully controlled experiments and to the discovery of deep laws of great generality is to study synthetic life forms in a synthetic world to which they have become coupled through evolution. This is the approach of *synthetic ethology*. Some simple synthetic ethology experiments are described in which we have observed the evolution of communication in a population of simple machines. We show that

**Artificial Life II: The Second Workshop on the Synthesis and Simulation of Living Systems*, Santa Fe Institute Studies in the Sciences of Complexity, proceedings Vol. X, edited by Christopher G. Langton, Charles Taylor, J. Doyne Farmer, and Steen Rasmussen. Redwood City, CA: Addison-Wesley, 1992, pp. 631-658.

even in these simple worlds we find some of the richness and complexity found in natural communication.

I am an “old bird,” ... a Simorg, an “all-knowing Bird of Ages” ...

— DeMorgan, *Budget of Paradoxes*, 1872, p. 329.

1 The Problem

Language, communication and other mental phenomena have been studied for many centuries, yet some of the central issues remain unresolved. These include the mechanisms by which language and communication emerge, the physical embodiment of mental states, and the nature of intentionality. I will argue below that answering these questions requires a deep theoretical understanding of communication in terms of the relation between its mechanism and its role in the evolution of the communicators. This is one of the goals of *ethology*, which “is distinguished from other approaches to the study of behaviour in seeking to combine functional and causal types of explanation.” [23] Our approach differs from traditional ethological methods in that it seeks experimental simplicity and control by studying synthetic organisms in synthetic environments, rather than natural organisms in natural environments; it is thus called *synthetic ethology*.

To explain why we expect synthetic ethology to succeed where other methods have failed, it is necessary to briefly review the previous approaches. In doing this I will focus on a single issue: How can a symbol come to *mean* something?

1.1 Philosophical Approaches

Although philosophical methods are quite different from those proposed here, the investigations of several philosophers lend support to synthetic ethology. To see why, consider the denotational theory of meaning, in which the meaning of a word is the thing that it denotes. This theory, which is commonly taken for granted, works well for proper names (‘Bertrand Russell’ denotes a particular person; ‘Santa

Fe' denotes a particular city), but becomes less satisfactory with increasingly abstract terms. Even for concrete general terms ('dog', 'mountain') it is already difficult to say exactly what they denote, as evidenced by 2500 years of debate over the nature of universals. Verbs are even more problematic, and a denotational theory of terms such as 'of' and 'the' seems hopeless.

In this century denotational theories of meaning came under attack from Wittgenstein and other "ordinary language" philosophers.[33] They pointed out that only a small number of linguistic forms can be understood in terms of their denotation; a more generally applicable theory must ground the meaning of language in its use in a social context. For example, in a simple question such as 'Is there water in the refrigerator?', the term 'water' cannot be taken to have a simple denotational meaning (such as a certain minimum number of H₂O molecules). Rather, there is a common basis of understanding, grounded in the speaker's and hearer's mutual interests and in the context of the utterance, that governs the quantity, state, purity, spatial configuration, etc. that a substance in the refrigerator should have to elicit a truthful "yes" response. To understand the meaning of 'water' we must know the function of the word in its contexts of use. Even scientific terms (e.g., length, mass, energy) acquire their meaning through measurement practices that form a common basis of understanding among scientists.

Heidegger makes very similar points, although with a different purpose.[12, 14, 15] He shows how our everyday use of language is part of a culturally constituted nexus of needs, concerns and skillful behavior. In his terms this nexus is a "world," and thus our linguistic behavior both is defined by and contributes to defining the various "worlds" in which we dwell: consider common expressions such as "the world of politics," "the academic world," and "the world of science." Meaning emerges from a shared cultural background of beliefs, practices, expectations and concerns. (Related ideas are discussed by Preston.[25])

One consequence of these views of language is that the study of language cannot be separated from the study of its cultural matrix. Thus one might despair that we will ever have a scientific theory of meaning. Fortunately, another philosopher, Popper, has shown a possible way out of this difficulty: "The main task of the theory of human knowledge is to understand it as continuous with animal knowl-

edge; and to understand also its discontinuity — if any — from animal knowledge.”[24] This is a very unconventional view of epistemology; traditionally philosophers have limited their attention to human knowledge, and in particular to its embodiment in human language. Although Heidegger and others have helped to bring non-verbal knowledge into the scope of philosophical investigation, Popper goes a step further, by indicating the importance of animal knowledge.

The importance of Popper’s observation for the study of language and the mind is that it encourages us to study these phenomena in the context of simple animals in simple environments. Science usually progresses fastest when it is able to study phenomena in their simplest contexts. We expect this will also be the case with communication and other mental phenomena: we will learn more if we start by studying their simplest manifestations, rather than their most complex (i.e. in humans).

1.2 The Behaviorist Approach

The preceding observations might suggest a *behaviorist* approach, since communication is a behavior and behaviorist experiments often involve simple animals in simple environments. But the behaviorist approach is inadequate for several reasons. First, it suffers from *ecological invalidity*. Animals behave in abnormal ways when put in alien environments, but what could be more alien than a Skinner box? As a result, the behavior of animals in laboratory situations may do little to inform us of their behavior in their natural environments.

Second, behaviorism investigates little snippets of behavior, such as pressing a lever to get some food, an approach that removes these behaviors from the pragmatic context that gives them their meaning. The result is an investigation of meaningless behavior resulting from a *lack of pragmatic context*. An example will illustrate the pitfalls of this approach. On the basis of behavioristic tests it had been thought that honey-bees were color-blind. However, von Frisch showed that in a feeding context they were able to distinguish colors. In the captive, laboratory context the color of lights was not relevant to the bees.[23]

In principle, of course, we could design experimental situations that mimic the natural environment in just the relevant ways and simplify it in ways that don’t distort the phenomena. Unfortunately, we don’t yet adequately understand the pragmatics of real life, and

so we don't know how to design laboratory environments that match the natural environments in just the relevant ways. Therefore, the behaviorist approach is, at very least, premature.

1.3 The Ethological Approach

An alternative approach to the study of communication is found in *ethology*, which is in part a reaction against behaviorism. Ethology recognizes that the behavior of an organism is intimately coupled (through natural selection) with its environment. Therefore, since removing an organism from its environment destroys the context for its behavior, ethology advocates studying animals in their own worlds (or in laboratory situations which closely approximate the natural environment). Unfortunately there are difficulties with this approach.

First, the real world, especially out in the field, is very messy; there are too many variables for clean experimental design. Consider some of the factors that could plausibly affect the behavior of a group of animals: the distribution of other animals and their behavior, the distribution of plants and their growth, the terrain, the weather, ambient sounds and odors, disease agents, etc. etc.[28] Animals are much too sensitive to their environments to permit a cavalier disregard for any of these factors.

Second, there are practical and ethical limits to the experiments we can perform. The ethical limits are most apparent where human behavior is the subject, but the situation differs only in degree where other animals are concerned. Even in the absence of ethical constraints, control of many variables is difficult.[28] Some of the experiments we would most like to perform are completely beyond our capabilities, such as restarting evolution and watching or manipulating its progress.

These two problems — the large number of variables and our inability to control them — make it unlikely that deep ethological laws will be discovered in the field. The history of the other sciences shows that deep, universal laws are most likely to be found when the relevant variables are known and under experimental control. When this is not the case, the best we can hope for is statistical correlation; causal understanding will elude us. Of course, I'm not claiming that empirical ethology is futile, only that it is very hard. Rather I anticipate that synthetic and empirical ethology are complementary approaches to the

study of behavior, and that there will be a fruitful exchange between them.

1.4 The Neuropsychological Approach

Behaviorist and ethological investigations of communication are limited in an additional way: they tell us nothing of the *mechanism* by which animals communicate. They are both based on *black-box* descriptions of behavior. On the other hand, deep scientific laws are generally based on a causal understanding of the phenomena. Thus it is important to understand the mechanism underlying meaning and other mental phenomena.[20] Several disciplines investigate the mechanisms of cognition. One is neuropsychology. Unfortunately, the complexity of biological nervous systems is so great that the discovery of deep laws seems unlikely, at least at the current stage of the science. Furthermore, as we've seen, true understanding of communication and other mental phenomena requires them to be understood in their ecological context. Thus a complete theory of communication must unite the neuropsychological and ethological levels. This is far beyond the reach of contemporary science.

1.5 The Artificial Intelligence Approach

Another discipline that investigates cognitive mechanisms is artificial intelligence, but with the goal of creating them, rather than studying their naturally occurring forms. Since AI creates its subject matter, all the variables are in its control, and so it might seem that AI is an ideal vehicle for studying communication, meaning and the mind. Unfortunately, as is well known, there's much argument about whether AI systems can — even in principle — exhibit genuine understanding. In other words, it is claimed that since AI systems perform meaningless (syntactic) symbol manipulation, they lack just the properties we want to study: meaningful (semantic and pragmatic) symbol use and genuine intentionality. I will briefly review the key points.

The issue can be put this way: Are AI programs *really* intelligent or do they merely *simulate* real intelligence. Several well-known examples make the difference clear. It has been pointed out that no one gets wet when a meteorologist simulates a hurricane in a computer; there is an obvious difference between a real hurricane and a simu-

lated hurricane.[26, 27] Similarly, it is been observed that thinking, like digestion, is tied to its biological context. The same chemical reactions will not be digestion if they take place in a flask, that is, out of the context of a stomach serving its functional role in the life of an organism. By analogy it is claimed that there cannot be any *real* thinking outside of its biological context: just as the flask is not digesting, so the computer is not thinking. It has also been claimed that computers may be able to *simulate* meaningful symbolic activity, but that symbols cannot *really* mean anything to a computer. In particular, any meaning born by machine-processed language is meaning that is derived from *our* use of the language. The rules we put into the machine reflect the meaning of the symbols to us; they have no meaning to the machine. That is, our linguistic behavior has *original intentionality*; whereas machines' linguistic behavior has only *derived intentionality*. [10, 11]

1.6 Summary

Here is the problem in a nutshell. If we want to understand what makes symbols meaningful (and related phenomena such as intentionality), then AI — at least as currently pursued — will not do. If we want genuine meaning and original intentionality, then communication must have real relevance to the communicators. Furthermore, if we are to understand the pragmatic context of the communication and preserve ecological validity, then it must occur in the communicators' natural environment, that to which they have become coupled through natural selection. Unfortunately, the natural environments of biological organisms are too complicated for carefully controlled experiments.

2 Synthetic Ethology as a Solution

2.1 Definition of Synthetic Ethology

The goal of *synthetic ethology* is to integrate mechanistic and ethological accounts of behavior by combining the simplicity and control of behaviorist methods with the ecological and pragmatic validity of empirical ethology. The idea of synthetic ethology is simple: Instead of studying animals in the messy natural world, and instead of ripping

animals out of their worlds altogether, we create artificial worlds and simulated organisms (*simorgs*¹) whose behavior is coupled to those worlds. Since the simulated organisms are simple, we can study mental phenomena in situations in which the mechanism is *transparent*. In brief, instead of *analyzing* the natural world, we *synthesize* an artificial world more amenable to scientific investigation. This is really just the standard method of experimental science.

Synthetic ethology can be considered an extension of Braitenberg's *synthetic psychology*[4] that preserves ecological validity and pragmatic context by requiring that behavior be coupled to the environment. We ensure this coupling by having the simorgs evolve in their artificial world. Synthetic ethology is also related to *computational neuroethology*,[1, 2, 8] the principal distinction being that that discipline typically studies the interaction of an *individual* organism with its environment, whereas our investigations require the study of groups of organisms.

2.2 Requirements of a Solution

In the following I argue that synthetic ethology does in fact solve the problems discussed above. First observe that, rather than starting with nature in all its glory, as does empirical ethology, or with denatured nature, as does behaviorism, synthetic ethology deals with complete, but simple worlds. Complexity is added only as necessary to produce the phenomena of interest, yet the worlds are complete, for they provide the complete environment in which the simorgs “live” or “die” (persist or cease to exist as structures).

Second, observe that because synthetic ethology creates the worlds it studies, every variable is under the control of the investigator. Further, the speed of the computer allows evolution to be observed across thousands of generations; we may create worlds, observe their evolution, and destroy them at will. Also, such use of simorgs is unlikely to be an ethical issue, at least so long as they are structurally simple.

Finally I claim that synthetic ethology investigates real, not simulated, communication. But how can we ensure that linguistic structures really “mean” something, that *communication* is taking place,

¹The simorg (simurg, simurgh), a monstrous bird of Persian legend, was believed to be of great age and capable of rational thought and speech.

and not merely the generation and recognition of meaningless symbols? Wittgenstein has shown that we are unlikely to find necessary and sufficient conditions governing our everyday use of words such as ‘communication’, and he has warned us of the pitfalls of removing words from their everyday contexts. Nevertheless, in the very non-everyday context of synthetic ethology, we need a definition that can be applied to novel situations.

As a first approximation, we might say that something is meaningful if it has relevance to the life of the individual. Perhaps we could go so far as to say it must be relevant to its survival — even if only indirectly or potentially. Relevance to the individual cannot be the whole story, however, since there are many examples of communication that do not benefit the communicator (e.g., the prairie dog’s warning call, a mother bird’s feigning injury). Thus, as a second approximation we can say that something is meaningful if it is relevant to the survival of the language community.

Additional support for this criterion comes from ethology, which has had to grapple with the problem of defining communication.[5, 9, 28, 30] The means that animals use to communicate, both within and between species, are so varied that identifying an act as communication becomes problematic. One animal scratches the bark of a tree; later another animal notes the scratches and goes a different way. Was it a communication act? The first animal might have been marking its territory, which is a form of communication, or it might simply have been sharpening its claws, which is not.

On the one hand we might say that a communication act has occurred whenever the behavior of one animal influences the behavior of another, but this definition is useless, since it views almost every behavior as communication. On the other hand we might say that it is not a communication act unless the first animal *intended* to influence the other’s behavior, but this criterion requires us to be able to determine the intent of behaviors, which is very problematic. If it is questionable to attribute intent to a fly, it is reckless to attribute it to a simorg: we need a definition of communication that does not appeal to problematic ideas like “intent.”

A definition of communication that is very consistent with our approach has been proposed by Burghardt:[5, 6]

Communication is the phenomenon of one organism producing a signal that, when responded to by another organ-

ism, confers some advantage (or the statistical probability of it) to the signaler or his group.

This says that communication must be relevant — in an evolutionary sense — to the signaler. In addition it gives us an operational way of determining if a communication act has taken place: we can compare the fitness of a population in the two situations differing only in whether communication is permitted or suppressed. This is the sort of experiment that can be undertaken in synthetic ethology, but that is infeasible for empirical ethology.

2.3 Making Real Worlds Inside the Computer

The objection may still be made that any communication that might take place is at best simulated. After all, nothing that takes place in the computer is real, the argument goes; no one gets wet from a hurricane in a computer. To counter this objection I would like to suggest a different way of looking at computers. We are accustomed to thinking of computers as abstract symbol-manipulating machines, realizations of universal Turing machines. I want to suggest that we think of computers as programmable mass-energy manipulators. The point is that the state of the computer is embodied in the distribution of real matter and energy, and that this matter and energy is redistributed under the control of the program. In effect, the program defines the laws of nature that hold within the computer. Suppose a program defines laws that permit (real!) mass-energy structures to form, stabilize, reproduce and evolve in the computer. If these structures satisfy the formal conditions of life, then they are real life, not simulated life, since they are composed of real matter and energy. Thus the computer may be a real niche for real artificial life — not carbon-based, but electron-based.² Similarly, if through signaling processes these structures promote their own and their group's persistence, then it is real, not simulated, communication that is occurring.

3 Preliminary Experiments

To illustrate the method of synthetic ethology, I will describe several experiments that have been completed. The goal of these experiments

²There is no claim here, however, that the simorgs used in these experiments are alive.

was to demonstrate that genuine communication could evolve in an artificial world. A secondary goal was to accomplish this with the simplest procedure possible, so that the phenomena would be most exposed for observation.³

3.1 Setup

3.1.1 Environment

What are the minimum requirements on a world that will lead to the emergence of communication? First, it must permit some simorgs to “see” things that others cannot, otherwise there would be no advantage in communicating. For example, in the natural world the signaler may perceive something which is out of the range of the receiver’s senses, or the signaler may be communicating its own internal state, which is not directly accessible to the receiver. Second, the environment must provide a physical basis for communication: something which the signaler can alter and the alteration of which the receiver can detect. Finally, we want the environment to be as simple as possible, so that the phenomena are manifest.

The solution adopted in these experiments is to give each simorg a *local environment* that only it can “see.” The states of the local environments, which we call *situations*, are determined by a random process; therefore there is no way they can be predicted. This means that the only way one simorg can reliably predict another’s situation is if the second simorg communicates that information to the first. To provide a medium for potential communication there is also a shared *global environment* in which any simorg can make or sense a *symbol*. Any such symbol replaces the previous contents of the global environment; there can be only one symbol in the “air” at a time. See Figure 1 for the topology of the environment.

In these experiments the situations and symbols (local and global environment states) are just natural numbers representing uninterpreted elements of a finite discrete set. Since we are creating an artificial world, there is no need to equip it with familiar environmental

³Our experiment may be contrasted with that of Werner and Dyer, who also observed the evolution of communication, but in a more complicated synthetic world.[31] That such different experimental designs resulted in qualitatively similar observations is evidence that synthetic ethology can reveal general properties of communication.




Figure 1: Topology of the Environment

features such as temperature, water supply, food supply, etc. We can define the laws of this universe so that the simorgs will survive only if they interact correctly with the uninterpreted states of this artificial environment. Although these states have no interpretable “meaning,” they are not simply syntactic, since they are directly relevant to the continued persistence (“survival”) of the simorgs.

3.1.2 Simorgs

Next consider the simorgs; they should be as simple as possible, yet be capable of evolving or learning complex behaviors. Two simple machine models have the required characteristics, although there are certainly others; they are *finite state machines* (FSMs) and *artificial neural networks* (ANNs). Although ANNs are better models for a variety of reasons,[19] we used FSMs in the experiments described here. (See our progress report for some ANN-based experiments.[22])

Finite state machines get their name from their internal memory, which at any given time is in one of a finite number of states. In addition, an FSM may have a finite number of sensors and effectors, the states of which are also finite in number. The behavior of an FSM is defined by its *transition table*, which comprises a finite number of

discrete rules. For each sensor state s and each internal state i , the table defines an effector state e and a new internal state i' . The machines used in these experiment have only one internal (memory) state. In other words, they have no ability to remember; therefore their response is completely determined by the current stimulus (i.e., their own situation and the shared symbol). In effect, each machine is defined by a table mapping symbol/situation pairs into responses.

There are two kinds of responses, *emissions* and *actions*. The effect of an emission is to change the symbol in the global environment, hence a response that is an emission must specify the symbol to be emitted. Actions are what must be accomplished effectively for the simorg to survive. Since we are selecting for cooperation we consider a simorg's action effective only if it matches the situation of another simorg. Thus a response that is an action must specify a situation that it is trying to match.

In these experiments we placed an additional requirement on effective action, namely that the action match the situation *of the last emitter*. This increases the selective pressure in favor of communication. Although one may find analogs of this in the natural world (e.g., a predator signaling for appropriate aid in bringing down some prey), the essential point is that we are making an artificial world and so we can define the laws to suit the needs of our experiment.

3.1.3 Fitness

The principal goal of the selective criteria is that they lead to the emergence of communication — without being overly “rigged.” In these experiments the environment selects for cooperative activity that requires knowledge of something that cannot be directly perceived, namely another simorg's local environment. Specifically, whenever a simorg acts, its action is compared to the situation of the simorg that most recently emitted. If the two match, then we consider an effective action to have taken place, and both the emitter and actor are given a point of credit. Since several simorgs may respond to a given emitter, a successful emitter can in principle accumulate considerable credit. Each simorg is given an opportunity to respond ten times before all the local environments are changed randomly.⁴ This interval

⁴In these experiments the simorgs were serviced in a regular, cyclic fashion. This means that communications with one's nearest neighbors in one direction (say clockwise) are least

is called a *minor cycle*. Credit is accumulated over a *major cycle*, which comprises five minor cycles. The resulting total is considered the simorg’s “fitness” for that major cycle, since it measures the number of times the simorg cooperated successfully; it is the criterion by which simorgs are selected to breed or die.

3.1.4 The Birth and Death Cycle

At the end of each major cycle, one simorg is selected to die and two simorgs are selected to breed. This keeps the size of the population constant, which simplifies the simulation and the analysis. Of course, we want the most fit to be most likely to breed and the least fit to be most likely to die.

For reasons discussed later (Section 3.2.3), we use the fitness to determine the *probability* of breeding or dying. In these experiments we made the probability of breeding proportional to the fitness (credit accumulated over one major cycle):

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

where p_k is simorg k ’s probability of breeding, ϕ_k is its fitness, P is the population size, and $\alpha = P^{-1} \sum_{k=1}^P \phi_k$ is the average fitness of the population. (If $\alpha = 0$ we set $p_k = 1/P$.) The probability of dying cannot in general be inversely proportional to fitness. However, we can make it a monotonically decreasing first-degree polynomial of fitness:

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

where q_k is the probability of dying and β is the fitness of the most fit simorg. (If $\alpha = \beta$ we set $q_k = 1/P$.)

The offspring is derived from its parents by a simplified genetic process. Each simorg has two transition tables, its genotype and its phenotype. The genotypes of the parents are used to determine the genotype of their offspring by a process described below. In general the genotype defines a developmental process leading to the phenotype, and the phenotype determines the simorg’s behavior. In these experiments this process is trivial: the initial phenotype *is* the genotype.

likely to be disrupted by other emitters. This may be important in forming “communities” using the same “language” (code).

Further, if learning is disabled (see Section 3.1.5), then the phenotype remains identical to the genotype.

The genotype of a simorg is a transition table, which defines a response for every symbol/situation pair. Each response is represented, in these experiments, by a pair of numbers, the first of which is 0 or 1, indicating *act* or *emit*, and the second of which is the situation or symbol that goes with the action or emission. The genome itself is just a string containing all these pairs; thus each “gene” defines the response to a given stimulus.

The (unmutated) genotype of the offspring is derived from its parents’ genotypes by a process called *crossover*. For purposes of crossover we interpret the genetic string as a closed loop. Two crossover points θ and θ' are selected randomly, and a new genetic string is generated from those of the parents. That is, between θ and θ' the genes will be copied from one parent, and between θ' and θ from the other. Note that our crossover operation never “splits its genes;” it cannot break up a transition table entry. We have found that this leads to faster evolution since the genetic operations respect the structural units of the genetic string. With low probability (0.01 in these experiments) the genetic string may be mutated after crossover. This means that a randomly selected gene is completely replaced by a random allele (i.e., a pair of random numbers in the appropriate ranges).

3.1.5 Learning

In order to experiment with the effects of learning on the evolution of communication, we have implemented the simplest kind of “single case learning.” Specifically, whenever a simorg acts *ineffectively* we change its phenotype so that it *would have acted* effectively. That is, suppose that the global environment state is γ and the local environment state is λ , and that under this stimulus a simorg responds with action λ' , but that the situation of the last emitter is $\lambda'' \neq \lambda'$. Then we replace the (γ, λ) entry of the *phenotypic* transition table with the action λ'' . (Of course, learning alters the phenotype, not the genotype.) This is a very simple model of learning, and could easily lead to instability; nevertheless it produces interesting results (see Section 3.2).

3.1.6 Importance of Overlapping Generations

Because we are interested in the influence of learning on the evolution of communication, we have done some things differently from typical genetic algorithms.[13, 17] GAs typically replace the entire population each generation, with the fitness of the parents determining the frequency with which their offspring are represented in the new generation. In contrast, we replace one individual at a time, with fitness determining the probability of breeding and dying. The difference is significant, because the GA approach prevents the passage of “cultural” information from one generation to the next (through learning). In the current experiment this happens indirectly, since symbol/situation associations are learned through ineffective action. Future experiments may model more direct transmission by having the less successful simorgs imitate the behavior of the more successful. We expect that “cultural” phenomena will be central to understanding the interaction of learning and communication. (See also Belew[3].)

3.1.7 Measurements

How can we tell if communication is taking place? As noted previously (Section 2.2), Burghardt’s definition of communication suggests an operational approach to identifying communication: detect situations in which one simorg produces a signal, another responds to it, and the result is a likely increase in the fitness of the signaler or its group.

In our case, fitness is a direct measure of the number of times that an effective action resulted from a simorg’s response to the last emitter. Therefore, the average fitness of the population measures the advantage resulting from actions coincident with apparent communication. But how do we know that the advantage results from communication, and not other adaptations (as it may; see Section 3.2.3)?

I have claimed that synthetic ethology permits a degree of control not possible in natural ethology, and here is a perfect example. We may start two evolutionary simulations with the same population of random simorgs. In one we suppress communication by writing a random symbol into the global environment at every opportunity; in effect this raises the “noise level” to the point where communication is impossible. In the other simulation we do nothing to prevent communication. If true communication — as manifested in selective advantage — is taking place, then the fitness achieved by the two pop-

ulations should differ. In particular, the *rate* of fitness increase should be significantly greater when communication is not suppressed. This is the effect for which we must watch.

In these experiments we record several fitness parameters. The most important is α , the average fitness of the population (smoothed by a rectangular window of width 50). The second most important is β , the fitness of the most fit simorg at the end of each major cycle (similarly smoothed). The figures in this chapter show the evolution of α ; the evolution of β is qualitatively similar.[21]

I am proposing synthetic ethology as a new way to study communication. Therefore, if by the process just described we find that communication is taking place, then we must see what the simulation can tell us about it. At this stage in the research program we have addressed only the most basic questions: What are the meanings of symbols, and how do they acquire them?

To answer these questions we construct during the simulation a data structure called a *denotation matrix*. This has an entry for each symbol/situation pair, which is incremented whenever there is an apparent communication act involving that pair. If symbols are being used in a haphazard fashion, then all the pairs should occur with approximately the same frequency; the matrix should be quite uniform. On the other hand, if the symbols are being used in a very systematic way, then we should expect there to be one situation for each symbol, and vice versa.⁵ Each row and each column of the denotation matrix should have a single nonzero entry, and these should all be about equal; this is a very nonuniform matrix, which we will call the *ideal* denotation matrix. Thus systematic use of symbols can be detected (and quantified) by measuring the *variation* (or *dispersion*) of the denotation matrix.

One of the simplest measures of variation is the standard deviation, which is zero for a uniform distribution, and increases as the distribution spreads around the mean. However the standard deviation is not convenient for comparing the uniformity of denotation matrices between simulations, since the mean may vary from run to run. Instead, we use the *coefficient of variation* (V), which measures

⁵This is assuming that the number of local environment states equals the number of global environment states, as it does in these experiments. We discuss later (Section 4) the consequences of having unequal numbers of states.

the standard deviation (σ) in units of the mean (μ):

$$V = \sigma/\mu.$$

The coefficient of variation is 0 for a uniform denotation matrix, and for the ideal matrix is $\sqrt{N-1}$ (where N is the number of global or local environment states, which are assumed equal).

Another measure of uniformity is the entropy of a distribution, which is defined:⁶

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

This is maximized by the uniform distribution; since there are N^2 equally likely states, its entropy is $2 \log N$. The minimum entropy $H = 0$ is achieved by the “delta distribution” (which makes all the probabilities zero except one). This is not so interesting, however, as the entropy of the ideal matrix, which is easily calculated to be $\log N$. To allow comparisons between simulation runs, we also use a “disorder measure”:

$$\eta = \frac{H}{\log N} - 1.$$

This is a scaled and translated entropy, which has the value 1 for a uniform matrix, 0 for the ideal matrix, and -1 for the “overstructured” delta matrix.

There are a variety of other statistical measures that may be used to quantify the structure of the denotation matrix. For example, χ^2 will be 0 for the uniform matrix and maximum for the ideal matrix. Fortunately the results we have observed so far are robust in that they are qualitatively the same no matter what statistics are used.

3.2 Results

Unless otherwise specified, the experiments described here used a population size $P = 100$ of finite state machines with 1 internal state. Since the number of local and global environment states were the same, $N = 8$, each machine was defined by a transition table containing 64 stimulus/response rules. Simulations were generally run for 5000 major cycles (one birth per major cycle).

⁶We use logarithms to the base 2, so that our entropy measure is more easily interpreted.




Figure 2: Average Fitness, Communication Suppressed and Learning Disabled

3.2.1 Effect of Communication on Fitness

Figure 2 shows the evolution of the (smoothed) average fitness (α) of a typical random initial population when communication has been suppressed and learning has been disabled. It can be observed to have wandered around the fitness expected for machines that are guessing, $\alpha = 6.25$. (The analysis may be found in an earlier report.[21]) Linear regression detects a slight upward trend ($\hat{\alpha} = 1.55 \times 10^{-5}$). This is a stable phenomenon across simulations, and is explained later (Section 3.2.3).

Figure 3 shows the evolution of the average fitness for the same initial population as Figure 2, but with communication permitted




Figure 3: Average Fitness, Communication Permitted and Learning Disabled




Figure 4: Average Fitness, Communication Permitted and Learning Enabled

(learning still disabled). Within 5000 major cycles the average fitness reaches $\alpha = 11.5$, which is significantly above the guessing level ($\alpha = 6.25$). Furthermore, linear regression shows that the average fitness is increasing over 50 times as fast as when communication was suppressed ($\dot{\alpha} = 8.25 \times 10^{-4}$ vs. $\dot{\alpha} = 1.55 \times 10^{-5}$). We conclude that in this experiment communication has a remarkable selective advantage.

Figure 4 shows the evolution of α for the same initial population, but with communication permitted and learning enabled. First observe that the average fitness begins at a much higher level ($\alpha \approx 45$) than in the previous two experiments. This is because each simorg gets ten opportunities to respond to a given configuration of local environment states. Since learning changes the behavior of a simorg so

Table 1: Average Measurements over Several Random Populations

Measurement	Comm/Learning		
	N/N	Y/N	Y/Y
α	6.31	11.63	59.65
$\dot{\alpha} \times 10^4$	0.36	11.0	28.77
V	0.47	2.58	2.65
H	5.81	3.79	3.87
η	0.94	0.26	0.29

that its response would have been correct, an incorrect response could be followed by up to nine correct responses (provided no intervening emissions change the global environment). The combination of communication and learning allowed the average fitness to reach 55, which is nearly five times the level reached without learning and nearly nine times that achieved without communication. The rate of fitness increase was $\dot{\alpha} = 2.31 \times 10^{-3}$, which is almost three times as large as that without learning, and nearly 150 times as large as that without communication.

We have observed quantitatively similar results in many experiments. Table 1 (adapted from an earlier report[21]) shows average measurements from several experiments that differ only in initial population.

To better understand the asymptotic behavior of the evolutionary process, we have run several simulations for ten times as long as those previously described. Figure 5 shows the evolution with communication permitted but learning disabled, and Figure 6 shows the evolution of the same initial population, but with communication permitted and learning enabled. In the first case average fitness reached a level of approximately 20.⁷ In the second (learning permitted) α seems to have reached an equilibrium value ($\bar{\alpha} = 56.6$ in fact); we can also observe an apparent “genetic catastrophe” at about $t = 45000$.

The greatly increased fitness that results from not suppressing the signaling process supports the claim that we are observing *genuine* communication. The communication acts have real relevance to the

⁷Under reasonable assumptions the maximum α achievable without learning by a homogeneous population can be calculated to be 87.5; details are presented elsewhere.[21]




Figure 5: Average Fitness, Communication Permitted and Learning Disabled




Figure 6: Average Fitness, Communication Permitted and Learning Enabled

Table 2: Denotation Matrix, Communication Suppressed and Learning Disabled

symbol	situation							
	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

$$V = 0.409451$$

$$H = 5.868233$$

$$\eta = 0.9560777$$

simorgs because they significantly affect the survival of the signaler and its group (cf. Burghardt’s definition, Section 2.2).

3.2.2 Analysis of Denotation Matrices

If genuine communication is taking place, then we ought to be able to observe it in more structured use of symbols; therefore we consider the structure of the resulting denotation matrices. First consider Table 2; this is the denotation matrix from the same simulation shown in Figure 2. In the absence of communication and learning we see a very uniform matrix, as measured by its coefficient of variation $V = 0.41$ and entropy $H = 5.87$, which is nearly the maximum possible, 6. This is also reflected in the disorder parameter $\eta = 0.96$; recall that a uniform matrix has $\eta = 1$ and an “ideal” matrix has $\eta = 0$.

Table 3 shows the denotation matrix that results when communication is permitted; even to the eye it is much more structured than Table 2. This is confirmed by our measurements: $V = 2.27$ (cf. $V = 2.65$ for the ideal matrix), $H = 3.92$, $\eta = 0.31$.

Finally, Table 4 is the denotation matrix resulting from both communication and learning. Qualitatively and quantitatively it is very

Table 3: Denotation Matrix, Communication Permitted and Learning Disabled

symbol	situation							
	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

$$V = 2.272352$$

$$H = 3.915812$$

$$\eta = 0.3052707$$

Table 4: Denotation Matrix, Communication Permitted and Learning Enabled

symbol	situation							
	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

$$V = 2.165397$$

$$H = 4.208782$$

$$\eta = 0.4029273$$

similar to Table 3, but slightly *less* structured. This phenomenon is even more apparent in longer simulations, such as the $t = 50000$ simulations shown in Figures 5 and 6. In these, evolution in the absence of learning produced a denotation matrix having $\eta = -0.2$, indicating an *overstructured* language, whereas evolution with learning produced a less structured language ($\eta = +0.2$) but higher fitness.⁸ This phenomenon seems to be consistent with research indicating that there is an optimal degree of structure[18] and that that optimum is more easily achieved with learning.[16]

The “ideal” denotation matrix has one symbol for one situation and vice versa; this is a structure that we might expect to see emerging. For example, in the denotation matrix in Table 3 there is at least one symbol that predominantly denotes a single situation: in 86% of its recent uses, symbol 4 denoted situation 6, in the remainder situation 7. Since these are the only two uses of symbol 4, it seems likely that the denotation matrix reflects two subpopulations (of unequal size) using the same symbol for different situations. More nearly equal subpopulations may be indicated by symbols such as 7, which is used for situations 1 and 3 with nearly equal frequency.

Symbols being used to denote several situations may also result from their being used equivocally by a single population; they could reflect an intermediate stage in the evolution to univocal symbol use. It is difficult to discriminate between these two possibilities on the basis of just the denotation matrix. Doing so requires more detailed analysis of the simorgs in the final population, a process which is straight-forward in synthetic ethology, since we have complete access to the structure of the simorgs. (Simple examples of this kind of analysis are presented in our report.[21])

The natural way to interpret the denotation matrix is by rows, which reflects the significance of a symbol to a recipient; ethologists sometimes call this the *meaning* of a signal.[7, 28, 29, 30] We can also look at the denotation matrix by columns, which shows the situation a signaler was expressing by a symbol; ethologists call this the symbol’s *message*.[7, 28, 29, 30] Sometimes the two are symmetric. For example, in Table 3 the meaning of symbol 4 is usually (86%) situation 6, and the message ‘situation 6’ is usually (61%) represented by symbol 4.

⁸See Tables 17 and 18 in our earlier report[21] for the denotation matrices resulting from these experiments.

On the other hand, asymmetries may occur. Symbol 6 usually (51%) means situation 0, but situation 0 is usually (62%) represented by symbol 1. Conversely, situation 2 is usually (98%) represented by symbol 1, but symbol 1 usually (71%) means situation 0.

Even in a synthetic ethology experiment as simple as this, we may begin to observe some of the richness and complexity of real communication. For example, in the actual “language” or code reflected in the evolved denotation matrix — as opposed to the ideal matrix given by theory — we find that there is rarely a one-to-one (univocal) correspondence between symbols and situations. Indeed, it is quite possible that a simorg will attach different significance to a symbol when it is received or when it is emitted; that is, a simorg need not associate the same meaning and message to a given symbol. If this is the case for simorgs, then it would seem foolish to assume that in human languages an utterance has the same pragmatic significance when it is spoken as when it is heard.

The denotation matrix captures the actual use of the code by the entire population over the last 50 major cycles of the simulation. In this sense it is an irreducible description of the message and meaning associated with every symbol. It is irreducible because any attempt to ignore the lesser entries and specify a unique denotational meaning for a symbol will misrepresent the facts of communication. In fact symbol 4 means situation 7 some (16%) of the time; this is part of the overall meaning of symbol 4 in *that* population at *that* time. To say that symbol 4 *really* means situation 6, and that the rest is noise, is a misrepresentation of the “language.”

Given that the denotation matrix is the irreducible description of the code, we see that the evolution of the code is mirrored in the evolution of the denotation matrix. Indeed, in the denotation matrix we may see the code as an emergent nonequilibrium system, which organizes itself by promoting the fitness of simorgs that behave in accord with its emerging structure.[20] This emerging structure is measured by the decreasing entropy of the denotation matrix.

Over time we may observe a changing constellation of meanings associated with a given symbol, and of the symbols representing a given message. We have already seen that these experiments indicate both synonymous and equivocal symbols. The experiments also exhibit both context-sensitive emission and context-sensitive interpretation of symbols. This is because the emission of a symbol by a

simorg may depend on the global environment (providing a context) as well as its local environment. Similarly, the response of a simorg to a symbol depends on its situation, which supplies a context. Finally, we observe that the differing use of symbols in various contexts makes it quite possible for every simorg to be using a different dialect of the “language” manifest in the denotation matrix. Even in these simple experiments we can begin to appreciate the complexity of the relation between symbols and their significance.

3.2.3 Other Observations

In the course of these experiments we have made several observations that provide some insight into the evolution of communication.

All of our experiments in which communication (and learning) is suppressed show a slight upward trend in fitness (see Figure 2 and Table 1). This is surprising, since in the absence of communication it would seem that there is no way to improve on guessing. However, that is not the case, and the way that it can occur is an interesting demonstration of the force of the evolutionary process. To see this, observe that our definition of effective action (Section 3.1.3) permits a kind of “pseudo-cooperation through coadaptation.” Specifically, a simorg is credited whenever its action matches the situation of the last emitter, which is also credited. Therefore, if the population contains a group of simorgs that emit only when they are in a fixed subset E of situations, then the possible states of the last emitter will not be equally likely; specifically states in E will be more likely than the other states. Under these conditions a simorg can “beat the odds” by always guessing a situation in E . The coadaptation of such “pseudo-cooperating” groups of simorgs seems to account for the increase of fitness even when communication is suppressed.

We checked this hypothesis in several ways. First, we compared simulations with the usual scoring algorithms to those in which fitness was credited by a match to *any* other simorg (vice just the last emitter); this eliminated the possibility of pseudo-cooperation. As expected, there was no trend in the average fitness. Second, we inspected the denotation matrices; doing so showed that emissions occurred in only a subset of the situations. Third, we calculated the expected average fitness for homogeneous populations and subsets E of the observed size. With the parameters we used, and the observed size 3

for E , we calculated the expected fitness to be $\alpha = 20.83$; in four simulations we observed $\alpha = 20, 29, 21, 23$. Together these are strong evidence in favor of the hypothesis.

Pseudo-cooperation can be eliminated by not favoring a match to the most recent emitter. Unfortunately, this removes much of the selective pressure toward communication (since it makes guessing almost as good a strategy as communicating) and therefore slows the simulations. For this reason we have retained the original scoring rule; in most cases pseudo-cooperation is a low level effect that is unintrusive and can be ignored.

Another observation arose from earlier, unsuccessful experiments. Recall that fitness determines the probability of breeding or dying; there is always a chance that the least fit will breed and that the most fit will die. In earlier experiments we used a simpler approach: breed the two most fit simorgs and replace the least fit. Thus the current algorithm is stochastic, whereas the older one was deterministic (except in the case of fitness ties). The change was made because we never observed the evolution of communication in the deterministic situation.

The reason seems to be as follows. Since only the two most fit simorgs breed, other good, but not great, simorgs are forever excluded from contributing to the gene pool. Since language is hard to get started, it is to be expected that nascent communicators will not be as fit as guessers. Language communities will never evolve, unless they have some chance of breeding, and this seems to be prevented by the brittleness of the deterministic algorithm.

4 Conclusions

I have argued that a complete understanding of language, communication and the representational capabilities of mental states will require a theory that relates the mechanisms underlying cognition to the evolutionary process. I also argued that the complexity of natural organisms makes it unlikely that such an integrated theory can be found by empirical ethology. Therefore synthetic ethology has been proposed as a complementary research paradigm, since carefully controlled experiments and deep theoretical laws are more likely to be achievable in the comparative simplicity of synthetic worlds.

As an example of synthetic ethology I have described experiments in which we have observed the evolution of communication in a population of simple machines. This was accomplished by constructing a world in which there is selection for cooperation and in which effective cooperation requires communication. The control granted by synthetic ethology permitted us to observe the evolution of *the same population* in two worlds (one in which communication was suppressed, the other permitted), and thus to measure the evolutionary effect of communication. Further, synthetic ethology affords complete access to the structure of the simorgs, thus exposing the mechanisms underlying their communication.

We are hopeful that synthetic ethology will prove a fruitful method for investigating the relation between linguistic and mental structures and the world. The experiments described here are just a beginning, and there are many directions in which to proceed. For example, if the number of situations exceeds the number of symbols, then we would expect the simorgs to string symbols together into “sentences”; this has already been observed, but more experiments are needed to discover the syntax that will emerge and the factors affecting it.

It also seems likely that the complexity of language reflects the complexity of the world. Our experiments to date have used environments that are in one of a finite number of discrete, atomic situations, and the resulting “languages” have been similarly simple. This suggests that we equip our synthetic worlds with environments containing objects in various relationships; we expect this to lead to categories of symbols analogous to the parts of speech (nouns, adjectives, etc.).

To date our experiments have been based on finite, discrete sets of symbols and situations, but much of the natural world is characterized by continuous variation, and both human and animal communication make significant use of continuously variable parameters (loudness, pitch, rate etc.). Ethological studies[32] suggest that discreteness — so called “typical intensity” — will emerge to the extent that communication is noisy, an easy variable to control in synthetic ethology. We hope to address this issue in future experiments and thus identify the principles underlying the emergence of discrete symbolic processes.

5 Acknowledgements

Research reported herein was supported in part by Faculty Research Awards (1988, 1989) from the University of Tennessee, Knoxville. I'm also grateful to members of the Cognitive Science Laboratory and the Ethology Program at the University of Tennessee for much helpful advice, and to Beth Preston and the editors for comments on earlier drafts.

References

- [1] Beer, Randall D. *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*. San Diego, CA: Academic Press, 1990.
- [2] Beer, Randall D., Hillel J. Chiel, and Leon S. Sterling. "A Biological Perspective on Autonomous Agent Design." *Journal of Robotics and Autonomous Systems* **6**, 169–186; reprinted in: *New Architectures for Autonomous Agents*, edited by Pattie Maes. Cambridge, MA: MIT Press, in press.
- [3] Belew, Richard K. "Evolution, Learning and Culture: Computational Metaphors for Adaptive Search." *Complex Systems* **4** (1990).
- [4] Braitenberg, Valentino. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press, 1984.
- [5] Burghardt, Gordon M. "Defining 'Communication'." In *Communication by Chemical Signals*, edited by J. W. Johnston, Jr., D. G. Moulton and A. Turk. New York, NY: Appleton-Century-Crofts, 1970, 5–18.
- [6] Burghardt, Gordon M. "Ontogeny of Communication." In *How Animals Communicate*, edited by Thomas A. Sebeok. Bloomington, IN: Indiana University Press, 1977, 71–97.
- [7] Cherry, Colin. *On Human Communication*, second edition. Cambridge, MA: MIT Press, 1966, 171, 244.
- [8] Cliff, D. T. "Computational Neuroethology: A Provisional Manifesto." University of Sussex, Cognitive Science Research Paper CSRP 162, May 1990.

- [9] Dawkins, Richard. "Communication." In *The Oxford Companion to Animal Behavior*, edited by David McFarland. Oxford, UK: Oxford University Press, 1987, 78–91.
- [10] Dennett, Daniel. *The Intentional Stance*. Cambridge, MA: MIT Press, 1987.
- [11] Dennett, Daniel. "The Intentional Stance." *Behavioral and Brain Sciences* **11** (September 1988): 495–546.
- [12] Dreyfus, Hubert L. "Introduction." In *Husserl, Intentionality and Cognitive Science*, edited by Hubert L. Dreyfus with Harrison Hall. Cambridge, MA: MIT Press, 1982, 1–27.
- [13] Goldberg, D. E. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley, 1989.
- [14] Heidegger, Martin. *Being and Time*, translated by J. Macquarrie and E. Robinson. New York, NY: Harper and Row, 1962.
- [15] Heidegger, Martin. *The Basic Problems of Phenomenology*, translated by Albert Hofstadter. Bloomington, IN: Indiana University Press, 1982.
- [16] Hinton, G. E., and S. J. Nowlan. "How Learning Can Guide Evolution." *Complex Systems* **1** (1987): 495–501.
- [17] Holland, J. H. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press, 1975.
- [18] Langton, C. G. "Computation at the Edge of Chaos: Phase Transitions and Emergent Computation." *Physica D* **42** (1990): 12–37; reprinted in: *Emergent Computation*, edited by Stephanie Forrest. North-Holland, 1990, 12–37.
- [19] MacLennan, Bruce J. "Logic for the New AI." In *Aspects of Artificial Intelligence*, edited by James H. Fetzer. Dordrecht, Holland: Kluwer Academic Publishers, 1988, 163–192.
- [20] MacLennan, Bruce J. "Causes and Intentions." *Behavioral and Brain Sciences* **11** (September 1988): 519–520.
- [21] MacLennan, Bruce J. "Evolution of Communication in a Population of Simple Machines." University of Tennessee, Knoxville, Computer Science Department Technical Report CS-90-99, January 1990.

- [22] MacLennan, Bruce, Noel Jerke, Rick Stroud, and Marc Van-Heyningen. “Neural Network Models of Cognitive Processes: 1990 Progress Report,” University of Tennessee, Knoxville, Computer Science Department Technical Report, 1990.
- [23] McFarland, David. *The Oxford Companion to Animal Behavior*. Oxford, UK: Oxford University Press, 1987, 153.
- [24] Popper, Karl. “Campbell on the Evolutionary Theory of Knowledge.” In *Evolutionary Epistemology*, edited by Gerard Radnitzky and W. W. Bartley III. La Salle, IL: Open Court, 1987, 117; original version: “Replies to My Critics.” In *The Philosophy of Karl Popper*, edited by Paul A. Schilpp. La Salle, IL: Open Court, 1974, 1061.
- [25] Preston, Beth. “Heidegger and Artificial Intelligence.” unpublished, Center for the Philosophy of Science, University of Pittsburgh, 1989.
- [26] Searle, J. R. “Minds, Brains, and Programs.” *Behavioral and Brain Sciences* **3** (1980): 417–424, 450–457.
- [27] Searle, J. R. “Is the Brain’s Mind a Computer Program?” *Scientific American* **262** (January 1990): 26–31.
- [28] Slater, P. J. B. “The Study of Communication.” In *Animal Behavior Volume 2: Communication*, edited by T. R. Halliday and P. J. B. Slater. New York, NY: W. H. Freeman, 1983, 9–42.
- [29] Smith, W. J. “Message–Meaning Analyses.” In *Animal Communication*, edited by T. A. Sebeok. Bloomington, IN: Indiana University Press, 1968, 44–60.
- [30] Smith, W. J. *The Behavior of Communicating: An Ethological Approach*. Cambridge, MA: Harvard University Press, 1977, Ch. 1. For message/meaning analysis see p. 19.
- [31] Werner, Gregory M, and Michael G. Dyer. “Evolution of Communication in Artificial Organisms.” *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, edited by Doyne Farmer, Christopher Langton, Steen Rasmussen, and Chuck Taylor, proceedings vol. XII. Redwood City, CA: Addison-Wesley, 1991.
- [32] Wiley, R. H. “The Evolution of Communication: Information and Manipulation.” In *Animal Behavior Volume 2: Communication*,

- edited by T. R. Halliday and P. J. B. Slater. New York, NY: W. H. Freeman, 1983, 156–189. For typical intensity see 163–164.
- [33] Wittgenstein, Ludwig. *Philosophical Investigations*, third edition. New York, NY: Macmillan, 1958, section 19.