

The Emergence of a "Language" in an Evolving Population of Neural Networks

Angelo Cangelosi ^{1,2} Domenico Parisi ²

(1) Centre for Neural and Adaptive Systems,
University of Plymouth

(2) Institute of Psychology,
National Research Council, Rome

Abstract

The evolution of language implies the parallel evolution of an ability to respond appropriately to signals (language understanding) and an ability to produce the appropriate signals in the appropriate circumstances (language production). When linguistic signals are produced to inform other individuals, individuals that respond appropriately to these signals may increase their reproductive chances but it is less clear what is the reproductive advantage for the languages producers. We present simulations in which populations of neural networks living in an environment evolve a simple language with an informative function. Signals are produced to help other individuals to categorize edible and poisonous mushrooms in order to decide whether to approach or avoid encountered mushrooms. Language production, while not under direct evolutionary pressure, evolves as a by-product of the independently evolving perceptual ability to categorize mushrooms.

Keywords: Language evolution, Genetic algorithm, Artificial Life, Symbol grounding

The emergence of a "language" in an evolving population of neural networks

1. Modelling the evolution of language and communication

Human language has evolved and, therefore, it probably had some positive influence on the reproductive success of the individuals that exhibited it. The evolutionary origin of language and the selective pressures that may have originated it, however, are difficult to investigate because of the limited evidence available. Therefore, it may be useful to try to simulate on a computer different evolutionary scenarios in order to state more clearly the underlying hypotheses and to determine more objectively and in more detail their consequences.

In fact, interest in studying language's origin and evolution using computer simulations has increased considerably in the last few years. Some researchers have explored the evolution of language using models that describe language as a set of signal-meaning pairs (Steels, 1996; 1997; Oliphant and Batali, 1996; Di Paolo, 1997). These authors use populations of agents that play a communication game via the exchange of signals. At each time step a signal is selected according to a matrix that assigns a probability value to each signal in correspondence to each meaning. With this simulation approach it is possible to study the different conditions that allow the evolutionary emergence of shared vocabularies. Steels and Vogt (1997) have experimented with adaptive language games in pairs of physically embodied robot agents. The language game includes six steps: establishing contact with the other robot, identifying the communication topic, categorizing the surrounding world, speaker's encoding of the communicative signal, listener's decoding of the signal, feedback from listener to speaker. This robotic approach has the advantage of stressing the emergence of grounded perceptual categories and the development of a shared vocabulary to talk about the world. Other researchers have used simulations to examine important linguistic phenomena such as the evolutionary emergence of a critical period for language development (Hurford, 1991) and the interaction between evolution and learning in the emergence of language (Kirby & Hurford, 1997).

Other simulative models have addressed topics in the evolution of animal communication such as the reliability of communication signals in aggressive behavior (de Bourcier & Wheeler, 1997) and the evolution of signal diversity, e.g., in mating songs (Werner & Todd, 1997), with interesting results. For example, Werner and Todd point out that the evolution of high levels of song diversity is possible when sexual rather than natural selection is the mechanism responsible for evolution - a finding that could be related to the high level of diversity among human languages and to the similarities between sexual and cultural selection mechanisms.

Some simulations use neural networks to model organisms and genetic algorithms to model evolution. For example, Saunders and Pollack (1996) have used recurrent neural networks and the GNARL (Saunders, Angeline, & Pollack, 1994) evolutionary algorithm to study the evolution of continuous communicative systems, that is, the exchange of real valued signals in different input and output channels. The interaction protocol involves small groups of agents (2 or 3) that exchange signals about a source of food with the evolved continuous signals functioning as modulators of the agents' behavior. (For a general discussion of how to study language with neural networks in an Artificial Life perspective, cf. Parisi, 1997.)

In the present paper we describe some simulations on the evolutionary emergence of a very limited "language", made up of just two one-word utterances, in a population of simple organisms living in a simple environment. The behavior of each organism is controlled by a neural network and the evolution of the population of organisms is modeled using a standard genetic algorithm. In our simulations the "language" is genetically inherited rather than culturally transmitted and, moreover, the utterances have no internal structure and no syntax. Therefore, our "language" is closer to animal communication than to human language. However, some processes (and problems) such as selective reproduction, changes in population variability, drift, etc. appear to be shared by both biological and cultural evolution and, furthermore, the functional aspect of language we are interested in (e.g., as an aid to categorization) can be studied even in languages without syntax. In any case, for simplicity we will refer to our "language" as language, without quotation marks.

The evolution of a language in a population of individuals implies the parallel evolution of two distinct abilities: the ability to produce the appropriate signals in the appropriate circumstances and the ability to understand these signals, that is, to respond to each signal with the appropriate behavior. Each ability does not make much

evolutionary sense without the other and, in fact, every individual in a language-using population tends to possess both abilities. Let us consider how these two abilities might evolve. Imagine two different scenarios. In one scenario (Webb, 1994) the male of one animal species emits a particular signal. The female perceives the signal and uses the signal's properties (e.g., directional information about its source) to find the male. When male and female are close enough, they can mate. In this scenario both the male's behavior of producing the signal and the female's behavior of responding to the signal by approaching its source are under separate selective pressures. Males emitting the signal are more likely to have offspring than nonemitting males, and females that respond with the appropriate behavior are more likely to have offspring than nonresponding females. In fact Werner and Dyer (1991; 1994) have been able to evolve this type of language in a population of neural networks.

But consider another scenario. An individual perceives the location of a mushroom but is unable to perceive the mushroom's detailed perceptual properties (e.g., its shape and color) because the mushroom is too distant or for some other reason. Hence, the individual cannot recognize if the mushroom is edible or poisonous and it cannot decide whether to approach and eat the mushroom or to avoid it. A conspecific which is nearby, unlike the first individual, can perfectly perceive the mushroom's properties. This other individual emits a particular signal if the mushroom is edible and a different signal if the mushroom is poisonous. Based on both direct information about the mushroom's location and the signal emitted by the conspecific the first individual responds by approaching and eating the mushroom in the first case and by moving away in the second case.

In this scenario the behavior of understanding the signals by responding appropriately clearly is under selective pressure. Individuals that respond appropriately to signals will tend to live longer and to have more offspring than individuals that do not respond appropriately. On the other hand, the evolution of the behavior of producing the appropriate signals is more mysterious. What is the advantage of producing the signal to the individual that produces it? Why should an individual that produces the appropriate signals live longer and have more offspring than other individuals that fail to do so? Since language requires the parallel evolution of linguistic production and linguistic comprehension, how can language evolve when it has a purely informative function and therefore it is advantageous to the receiver but not to the producer?

2. Experimental setup: Living in an environment with both edible and poisonous mushrooms

In this Section we describe a scenario inspired by communicative signals in small groups of animals such as the well known signals with which vervet monkeys communicate the presence of various types of predators to their conspecifics (Cheney and Seyfarth, 1990). In the animal kingdom signals that refer to entities in the environment are used to refer not only to predators but to other entities as well. Use of signals to communicate information about food location and quality is present in many animal species (Hauser, 1996). Our scenario simulates the exchange of communicative signals between pairs of organisms concerning the quality of potential food. More specifically, individual organisms signal to each other if encountered mushrooms are edible or poisonous.

The organisms live in an environment that contains two types of mushrooms: edible mushrooms and poisonous mushrooms. Edible mushrooms resemble but are not identical to each other, and the same is true for poisonous mushrooms. Edible mushrooms are different from poisonous mushrooms but mushrooms belonging to the two different categories may share some properties. Since the organisms reproduce on the basis of their ability to eat the edible mushrooms and to avoid the poisonous ones, they must first categorize an encountered mushroom as either edible or poisonous and then they must respond by approaching and eating edible mushrooms and by going away from poisonous ones.

Each individual lives in an environment of $20 \times 20 = 400$ cells that contains 20 randomly distributed mushrooms each occupying a single cell. Ten mushrooms are edible and the other 10 are poisonous. At the beginning of its life an individual organism is placed in a randomly selected cell with a randomly selected orientation (N, S, E, and W). The entire life of an organism lasts 750 time units (input/output cycles) divided up into 15 "epochs" of 50 cycles each. When an organism happens to step on a cell containing a mushroom, the mushroom disappears (it is eaten). At the beginning of each epoch all the mushrooms remaining from the preceding epoch are eliminated and they are replaced by a new set of 20 randomly distributed mushrooms.

The behavior of each organism is controlled by a feedforward neural network with 14 input units, 5 output units, and 5 hidden units (Figure 1). One input unit encodes the location (direction) of the single nearest mushroom as the mushroom's angle measured clockwise from the organism's current facing direction. This angle is mapped in the interval from 0 to 1. (If two or more mushrooms are at the same distance from the organism, one is chosen at random.) Ten input units encode the mushroom's perceptual properties. The 10 edible mushrooms are encoded as 10 patterns of 10 bit, with each pattern obtained by changing a single bit, randomly chosen, in the prototypical pattern 1111100000. Similarly, the 10 poisonous mushrooms are encoded as 10 single-bit deviations from the prototype 0000011111. Hence, an edible and a poisonous mushroom can share either zero (when the two bit-changes neutralize each other) or two bits in the same position. The 3 remaining input units (signal-encoding input units) encode one of 8 possible perceived signals: 111, 110, 100, etc.

(Figure 1 about here)

Two of the 5 output units encode a movement of the organism in the environment. The organism can either proceed one step forward (11), turn 90 degrees to the left (10) or to the right (01), or just do nothing (00). The remaining 3 output units (signal-encoding output units) encode one of 8 possible emitted signals in the same way as the signal encoding input units. (For all output units continuous values are thresholded to either 0 or 1.)

We generate an initial population of 100 neural networks with the same architecture and randomly assigned connection weights. These weights are randomly chosen from a rectangular distribution ranging from -1 to +1. At the beginning of life an individual has zero energy. The individual's energy is increased by 10 units every time the organism eats an edible mushroom and it is decreased by 11 units if the organism eats a poisonous mushroom. At the end of life, which has the same length for all organisms, the organisms are ranked in terms of their energy and the 20 individuals with the most energy are allowed to reproduce by generating 5 offspring each. An offspring has the same connection weights of its (single) parent with the exception of some "genetic mutations" that change the value of 10% of the weights, randomly selected, by adding a quantity randomly chosen in the range -1/+1. The process is repeated for 1000 generations. The selective reproduction of the individuals with most energy and the constant addition of variation to the genetic pool of connection weights through the genetic mutations results in an increase in average energy across the 1000 generations and the evolutionary emergence of the behavior of approaching and eating the edible mushrooms and avoiding the poisonous ones.

However, we are not interested in this behavior as such. We are interested in how our organisms recognize edible and poisonous mushrooms and, more specifically, what we want to know is if a useful language will emerge in the population because it helps the organisms to discriminate between edible and poisonous mushrooms. Linguistic signals can be useful both as an aid in the categorization of nearby mushrooms whose perceptual properties are perceived by the organism and as "symbols" standing for the perceptual properties of more distant mushrooms that cannot be categorized as either edible or poisonous because while their location is perceived their perceptual properties are not perceived.

Imagine an organism wandering in its environment. In each cycle one particular mushroom happens to be the mushroom closest to the organism. If the mushroom is sufficiently near to the organism, i.e., it is located in one of the 8 cells adjacent to the organism's cell, the organism perceives both the location of the mushroom (its angle with respect to the organism's facing direction) and its perceptual properties (the pattern of 10 bits). However, if the mushroom is more distant, the organism can perceive the mushroom's location but not its perceptual properties. The 10 input units encoding the mushroom's perceptual properties all have 0 activation value.

We compare the evolution across 1000 generations of three different populations. One population has no language. In the organisms of this population the input units of the neural network that should encode the perceived signal are constantly kept at an activation level of 0.5 and the output of the signal-encoding output units is ignored. Like all the organisms in the present simulations, when an individual encounters a mushroom which is not located in one of the 8 cells adjacent to the individual's cell, the organism can perceive the direction in which the mushroom lies but not the mushroom's perceptual properties. Unlike the organisms of the next two populations, however, this individual is not helped in recognizing the type of mushroom by some linguistic signal. Therefore, the only solution which is open to this organism consists in approaching the unknown mushroom until it happens to be sufficiently near that the organism can perceive its perceptual properties. At this point the individual can categorize the mushroom and it can either eat the mushroom or go away from the mushroom as appropriate.

In a second type of population the language is externally provided by us and it does not evolve (Floreano, Miglino, & Parisi, 1991). When an individual belonging to this population encounters a mushroom, the three input units of its neural network that encode perceived signals have an activation pattern of '100' if the encountered mushroom is edible and an activation pattern of '010' if it is poisonous. (The first two digits discriminate between the two signals while the third digit has a value shared by both signals.) The signals produced by the organisms are ignored. These organisms can use the linguistic signals externally provided to them both as a help in categorizing a mushroom when they can perceive the mushroom's perceptual properties because the mushroom is close enough, and as a substitute for the act of perceiving the mushroom's perceptual properties when the encountered mushroom is too distant for the organism to have access to its perceptual properties.

In the third type of population language is not externally provided by us but it evolves autonomously. The scenario, which has been inspired by Hutchins and Hazelhurst (1995), is the following. Like the organisms of the other two populations, an individual can perceive the nearest mushroom's perceptual properties only if the mushroom is close enough. However, in this simulation something special happens. In each cycle another individual is randomly selected from the population (i.e., from the remaining 99 individuals forming the current generation) and this second individual is placed next to the first individual so that it is exposed to the same perceptual input as the first individual with the only difference that the second individual has access to the perceptual properties (the pattern of 10 bits) of the mushroom whatever the distance of the mushroom. The only task for the second individual is to label the mushroom for the first individual. The signal-encoding input units of the second individual are always set to 0.5. The output of its signal-encoding output units (thresholded to either 0 or 1) in response to the perceptual properties of the mushroom is used as input to the signal-encoding input units of the first individual. (The second individual's motor output is ignored.)

Therefore, in this last population when an individual encounters a mushroom the individual has always access to a linguistic signal produced by a conspecific. As in the previous population, if the mushroom is close enough so that not only the mushroom's location but also its perceptual properties can be perceived, the linguistic signal provided by the conspecific can be used by the individual as a help in categorizing the mushroom. If the mushroom is more distant and its perceptual properties are not accessible, the linguistic signal can function as a substitute for these perceptual properties. However, in this population, unlike the previous population, the quality of the signals provided by conspecifics is not guaranteed. Whatever signal is generated by the conspecific's neural network, the signal is input to the neural network of the individual that must decide whether to approach or go away from the mushroom. Hence, the language can be useful to these organisms only if it evolves appropriately.

3. Results

Figure 2 shows how average energy changes across 1000 generations in the three populations with no language, externally provided language, and evolved language, respectively (average results of 5 replications of each simulation). The simulations were stopped after 1000 generations because at that point the organisms are able to discriminate sufficiently well between edible and poisonous mushrooms and to associate the appropriate behavior to each type of mushrooms. A behavioral test of mushroom discrimination at generation 1000 has showed that in average a good organism collects 28 edible mushrooms and only 1 toadstool.

Language appears to be a useful addition to the evolutionary adaptation of these organisms. The organisms with no language have an average energy of a little more than 150 units at the end of evolution while the two populations with language have an average energy of more than 250 units. On the other hand, the two populations with language do not differ very much from each other. Although, predictably, the population with externally provided language has a more regular increase in average energy than the population with evolved language, the two populations reach an equivalent level of energy at the end of evolution.

(Figure 2 about here)

It is interesting to examine what linguistic signals evolve in the third population. Since the individuals that label the mushrooms are each time randomly selected from the population different individuals can produce different signals when exposed to the same perceptual input. What is observed, therefore, is a frequency distribution of different signals produced by the different individuals in each generation. To determine what signals are actually produced, each individual was tested using a sort of 'naming task' in an experimental controlled setting. The individual was exposed to the entire set of 20 mushrooms (10 edible and 10 poisonous ones) each positioned in four different locations (directions) with respect to the individual, that is, in front, in the back, right and left. The $20 \times 4 = 80$ signals produced by the individual in response to this set of 80 inputs were recorded and analyzed. (The signal-encoding input units of the individual were set to 0.5 during the entire test.)

Figure 3 shows the frequency distribution of the 8 possible signals produced by all the individuals in each of 10 generations (generation 0, 100, 200, etc., up to generation 1000) in one replication of the simulation. A similar pattern of frequency distribution was found in the other four replications. However the two high-frequency signals used to name the two types of mushrooms changed in each replication.

(Figure 3 about here)

Although there are some oscillations, the population evolves a language that tends to consistently use the pattern '010' to label edible mushrooms and the pattern '110' to label poisonous mushrooms. A population can be said to possess an efficient language if (a) functionally distinct categories (in our case, edible and poisonous mushrooms) are labeled with distinct signals, (b) a single signal tends to be used to label all the instances within a category, (c) all the individuals in the population tend to use the same signal to label the same category. (Clark (1993) has argued that principles similar to these govern the child's acquisition of the lexicon.) According to these criteria, the language evolved by our population appears to be rather efficient. (Similar results were obtained in the other replications of the simulation although of course different pairs of signals emerged for the two categories of mushrooms.)

4. Discussion

We have seen that a population of simple artificial organisms living in a simple environment can evolve an efficient language with an informative function to help the individuals to interact with their environment. Due to sensory limitations an individual can perceive the location but not the perceptual properties of a distant mushroom. This represents a serious handicap because an individual can adopt an informed decision on whether to approach or go away from an encountered mushroom only if the mushroom is very close. In these circumstances the population evolves a simple language in the sense that individuals tend to generate distinctive labels for edible and for poisonous mushrooms and these labels are used by other individuals to decide whether to approach or avoid a mushroom.

As we have observed in Section 1, the evolution of a language implies the parallel evolution of the ability to produce the appropriate signals in the appropriate circumstances and the ability to understand the perceived signals by responding appropriately to them. In some populations the two abilities co-evolve because there are separate evolutionary pressures on both language producers and language understanders. However, in the evolutionary scenario that we have studied in our simulations it is less clear how a language can evolve. When linguistic signals are used with the function to inform the receiver about some environmental object or event that is useful for the receiver to know more about, there is an evolutionary pressure on receivers of signals to evolve an ability to understand the signals. However, it is less clear what are the evolutionary pressures on the producers of the signals to evolve an ability to produce the appropriate signals in response to the appropriate object. But if the appropriate signals are not produced the ability to understand those signals does not make much sense and no language can evolve.

A possible answer to this apparent puzzle links the evolution of language to the evolving cognitive (perceptual) ability to categorize mushrooms (or, more generally, entities in the environment) based on their perceptual properties. If one examines the architecture of the neural network that governs the behavior of our organisms, one can say that the two abilities of producing and understanding linguistic signals are incorporated

(represented) in two separate sets of connection weights. The ability to understand the signals is represented in the lower connections weights linking the signal-encoding input units to the hidden units ("understanding weights"). The ability to produce signals is represented in the higher weights from the hidden units to the signal-encoding output units ("producing weights") (cf. Figure 1). The two sets of connections weights are interlinked through the layer of hidden units. So let us examine more closely the role of these hidden units.

In a feedforward neural network the input-to-hidden connection weights have the task to transform the input so that the input can be mapped into the appropriate output using the hidden-to-output connection weights. The result of this transformation is the activation pattern which is observed in the hidden units. In our networks the hidden units should be able to encode the transformed perceptual properties of encountered mushrooms in such a way that the output units can generate two different motor behaviors in response to the two categories of mushrooms: approaching and eating the edible mushrooms and going away from the poisonous ones. The hidden units can accomplish this task by exhibiting the same or very similar patterns of activation in response to all edible mushrooms and a distinct set of similar patterns in response to all poisonous mushrooms. At the beginning of evolution, given the random weights, the patterns evoked by the two categories of mushrooms are likely to overlap considerably. However, the evolutionary emergence of the ability to approach the edible mushrooms and to avoid the poisonous ones implies that the two sets of hidden activation patterns tend progressively to separate. The networks that tend to respond to all edible mushrooms with a given set of similar activation patterns in their hidden units and to poisonous mushrooms with a distinct (nonoverlapping) set of similar activation patterns are more likely to leave descendants.

Now let us introduce linguistic signals and let us consider how the understanding weights and the production weights are related to the activation patterns in the hidden units. Given some particular signal as input the role of the understanding weights is to help the hidden units to exhibit the appropriate activation pattern (tendentially a single one for all edible mushrooms and a different one for all poisonous mushrooms) when a mushroom is close enough and its perceptual properties can be perceived or, more critically, to induce the appropriate activation pattern in the hidden units when the mushroom is distant and its perceptual properties cannot be perceived. Hence, there is a direct evolutionary pressure to develop the ability to understand linguistic signals because understanding linguistic signals helps, or makes it possible, to categorize the different environmental entities and, as a consequence, to respond appropriately to these entities.

But it is the role of the hidden units in language production which is of interest here. This role may suggest how the evolutionary improvement in the cognitive (perceptual) ability to recognize different categories of mushrooms can have a positive influence on the evolution of the linguistic ability to produce an efficient language. Individuals that tend to exhibit one set of similar activation patterns on their hidden units in response to all edible mushrooms and a distinct set of similar activation pattern in response to all poisonous ones are more likely to designate with a distinct linguistic signal all edible mushrooms and with a different linguistic signal all poisonous mushrooms, that is, to produce a useful language. This results from the fact that the signals emitted by an individual in response to a perceived mushroom depend on two factors: (a) the individual's production weights, and (b) the activation pattern on the individual's hidden units (cf. Figure 1). In the early stages of evolution both factors are unable to produce useful signals. The production weights are initially assigned at random and they are not selected for producing useful signals. Furthermore, the activation patterns on the hidden units will vary inconsistently from one mushroom to another one. Hence, the activation pattern on the signal-encoding output units (i.e., the emitted signal) will tend to consist of values that vary more or less randomly around 0.5 (because the production weights have been randomly assigned at the beginning of evolution and random weights tend to result in 0.5 activation values) and, after these values have been thresholded to either 0 or 1, the resulting signal will not be very useful.

However, after a certain number of generations the organisms will be more able to categorize the encountered mushrooms appropriately with their hidden units because this is the only way for them to increase their reproductive chances. Better categorization, as we have seen, means that the hidden units tend to exhibit very similar activation patterns for all edible mushrooms and a distinct set of similar activation patterns for all poisonous mushrooms. Hence, factor (b) above has changed. More consistent activation patterns on the hidden units can result in slight changes in the activation level of the signal-encoding output units that may be sufficient to allow organisms to produce more useful signals, that is, a single signal for all edible mushrooms and a different signal for all poisonous mushrooms. Furthermore, the evolutionary increase in the absolute value of connection weights due to mutations can protect the production weights from the disrupting effects of mutations and consolidate this trend toward the production of better signals.

This analysis is confirmed by an examination of the signals produced (but not used) by the population without language. As will be recalled, the individuals in this population do not receive any linguistic signals in their signal-encoding input units (which have a constant activation value of 0.5) and the signals they produce are simply ignored. Therefore, there is no evolutionary pressure of any kind for the emergence of a language in this population. However, if we examine the activation patterns appearing on the signal-encoding output units of the individuals of our populations without language, i.e., the linguistic signals they produce even if nobody is using them, we see an interesting result. Figure 4 shows the evolution of average energy in one of our populations without language and at the same time it plots an index of the quality of the language produced based on the three criteria for an efficient language mentioned at the end of Section 3.

To compute the index of language quality we use an 8x2 language production table. The table contains the percentages of each of the 8 possible signals for the 2 classes of edible and poisonous mushrooms based on the results of the naming test for all individuals. The quality index QI is computed using the following formula:

$$QI = \sum_{i=1}^8 |x_i - y_i| - k * \min(d_{poisonous}, d_{edible}) \quad (1)$$

where x_i is the percentage of signal i for poisonous mushrooms and y_i for edible mushrooms; k (here 1) is a constant to weigh the effect of the internal dispersion value of poisonous or edible mushrooms. The dispersion values $d_{poisonous}$ and d_{edible} are computed with formula (2) using the expected percentage x_e and y_e in case of a flat distribution in which all 8 signals are equally used to name a class of mushroom:

$$d_{poisonous} = \sum_{i=1}^8 |x_i - x_e| \quad d_{edible} = \sum_{i=1}^8 |y_i - y_e| \quad (2)$$

While the first part of formula (1) measures the principle of contrast (use of only one word for each class of mushrooms), the dispersion values measure the use of synonyms for the same mushroom class. The use of the percentage table for the signals emitted by all organisms allows us to consider how much all the organisms share the same language.

Comparing the evolutionary trend of fitness and of the language quality index what is observed is that after a certain number of generations the organisms start producing useful linguistic signals and that, interestingly, the sudden improvement in the usefulness of their language coincides with a marked increase in their fitness (cf. Figure 4).

(Figure 4 about here)

The observed improvement in the quality index of the emitted, even if not used, language since generation 450 appears to be related to the sudden increase in average fitness which is also observed in this period. The increased fitness is at least in part determined by a better categorization of encountered mushrooms by the activation patterns in the hidden units. These better activation patterns of the hidden units in turn cause the production of better linguistic signals. Hence, the linguistic ability to produce useful signals appears to be a by-product of the cognitive ability to categorize perceived mushrooms. If the produced signals are received by other individuals, as in our third population, they may help these other individuals to categorize perceived mushrooms or they may function as substitutes for the perceptual properties of mushrooms when these properties cannot be perceived.

The fact that language production and categorization are related does not mean that any change in each of the two abilities should necessarily affect the other. In the simulations where language has no direct effect on the organism's fitness, the naming behavior depends also on the random drift of the connection weights between the hidden units and the output signal-emitting units. This probably explains why at about generation 950 the temporary decrease of language quality is not associated with a decrease in fitness.

In simulations where foraging behavior depends directly on language, the correlation between categorization ability and the quality of produced language can be much more important. We have measured this correlation using Pearson r . The correlation between average fitness and the language quality index for all five replications

is .71 ($p < .0001$). In this case any relevant change in the language or in the categorization skill will affect the other ability.

As we have already said, the positive effect of categorization on emitted language is explained in our model by their sharing the common layer of hidden units. Even though the present neural network architecture is not intended to reflect any specific animal or human brain structure, the shared layer of hidden units is assumed to correspond to the brain level in which perceptual and linguistic processing share some common information.

The influence of the perceptual ability to categorize entities in the environment on the linguistic ability to produce efficient signals to label these entities is related to a hypothesis about language evolution which has been proposed by Burling (1993). Burling has argued that human language has emerged from the cognitive (sensory-motor) capacities of our prelinguistic ancestors rather than from their primate-level communicative behavior. One might read our simulations as providing some evidence in favour of this hypothesis. The production of efficient linguistic signals as aids to categorization can be viewed as a by-product of, or at least to be facilitated by, a pre-existing ability to categorize the environment on a sensory-motor basis.

In our analysis of the language produced by the organisms in our third simulation (evolved language) we found that the evolved language tends to have the following three properties: (a) different signals are used for functionally distinct categories, (b) the same signal is used to label all the instances within a category, (c) all individuals use the same two signals for the two categories. The use of different signals for different categories and of the same signal for all members within a category corresponds to the principle of contrast (Clark, 1987) or of mutual exclusivity (Markman and Wachtel, 1988), that children rely on when they assign only one label per category. The fact that all individuals use the same signal for each category corresponds to the ontogenetic principle of conventionality (Clark, 1993). This correspondence between ontogenetic principles of language acquisition and phenomena of both language evolution and historical changes has been discussed by Clark (1993). For example, she compares the principle of contrast with Bréal's law of linguistic differentiation. In his work on historical changes of linguistic forms Bréal (1897) showed that a single meaning originally associated with several different signals becomes associated in later historical stages with only one of these signals.

A final aspect of our model that we would like to mention is its potential to deal with the symbol grounding problem (Harnad, 1990). Harnad has proposed the use of some kind of hybrid connectionist-symbolic system in which the connectionist component would generate symbols that are grounded in the physical world by perception and the symbolic component would manipulate these symbols (Harnad, 1993). The present model seems to be able to evolve signals that are linked to objects in the environment in that they are produced and understood by organisms as part of their sensory-motor interactions with the environment. The signals produced by our networks with their signal-encoding output units can be identified with Harnad's "symbolic representations", while the hidden unit activations correspond to "categorical representations" since their patterns tend to maximize inter-categorical differences and minimize intra-categorical differences. (For a discussion of these differences cf. Sharkey and Jackson, 1994; Cf. also Parisi, Denaro, and Cangelosi, 1997.) A further step in the evolution of a truly human language, of course, would be the evolution of an ability in neural networks to combine signals (and the internal categorical representations associated with them) to form complex signals with a syntax. If this can be done, our model would deal with signal manipulation and combination within a non-hybrid entirely connectionist framework.

5. Conclusions

The results of our efforts to simulate language evolution using neural networks in an ecological perspective have shown that the model proposed allows us to study at least some aspects of the emergence of language in populations of organisms interacting among themselves and with an external environment. The analysis of the relation between categorization and linguistic production suggested that there might be a strong interdependence between the evolution of language and the evolution of cognition and that language might have initially evolved as a by-product of cognitive (perceptual) abilities to categorize the environment. On the other hand, the role of language in the categorization of entities in the environment can have been one of the evolutionary pressures for the further evolution of categorization abilities.

In further development and testing of the model (Cangelosi, 1997; Cangelosi, Denaro, and Parisi, in preparation) we have investigated other aspects of the model such as the influence of language on categorization when

categorization is facilitated because individuals learn during their life (by imitating others) to label encountered mushrooms. In these simulations we have also directly inspected the internal representations of encountered mushrooms (activation patterns on hidden units) and compared internal representations without language and with perceived or produced language, and how these internal representations change during the process of evolution.

References

- Breal, M. (1897). *Essai de semantique*. Paris: Hachette.
- Burling, R. (1993). Primate calls, human language, and nonverbal communication. *Current Anthropology*, **34** (1), 25-53.
- Cangelosi, A. (1997). Towards new computational models of language: The ecological neural networks. *PhD Dissertation*, University of Genoa.
- Cangelosi, A., Denaro, D., & Parisi, D. (in preparation). An ecological neural network model of language learning and evolution.
- Cheney, D.L., & Seyfarth, R.M. (1990). *How monkeys see the world: Inside the mind of another specie*. Chicago, IL: Chicago University Press.
- Clark, E.V. (1987). The Principle of Contrast: a constraint on language acquisition. In B. MacWinney (Ed.), *Mechanisms of language acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Clark, E. (1993). *The lexicon in acquisition*. Cambridge, MA: Cambridge University Press.
- de Bourcier, P., & Wheeler, M. (1997). The truth is out there: The evolution of reliability in aggressive communication systems. In P. Husbands & I. Harvey (eds.), *Proceedings of the Fourth European Conference on Artificial Life*, Cambridge, MA: MIT Press.
- Di Paolo, E. (1997). Social coordination and spatial organisation: Steps toward the evolution of communication. In P. Husbands & I. Harvey (eds), *Proceedings of the Fourth European Conference on Artificial Life*, Cambridge, MA: MIT Press.
- Floreano, D., Miglino, O., & Parisi, D. (1991). Emerging complex behaviors in ecosystems of neural networks. In E.R. Caianiello (Ed.) *Parallel architectures and neural networks*. Singapore: World Scientific.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, **42**, 335-346.
- Harnad, S. (1993). Grounding symbols in the analog world with neural networks. *Think*, **2**, 12-78.
- Hauser, M.D. (1996). *The evolution of communication*. Cambridge, MA: MIT Press.
- Hurford, J.R. (1991). The evolution of the critical period for language acquisition. *Cognition*, **40**, 159-201.
- Hutchins, E., & Hazelhurst, B. (1995). How to invent a lexicon. The development of shared symbols in interaction. In N. Gilbert & R Conte (Eds.) *Artificial societies: The computer simulation of social life*. London: UCL Press.
- Kirby, S., & Hurford, J. (1997). Learning, culture and evolution in the origin of linguistic constraints. In P. Husbands & I. Harvey (eds), *Proceedings of the Fourth European Conference on Artificial Life*, Cambridge, MA: MIT Press.
- Markman, E.M., & Wachtel, G.F. (1988). Children's use of mutual exclusivity to constraint the meaning of words. *Cognitive Psychology*, **20**, 121-157.
- Oliphant, M., & Batali, J. (1996). Learning and the emergence of coordinated communication. *Technical Report*, University of California at San Diego, Department of Cognitive Science.
- Parisi, D. (1997). An Artificial Life approach to language. *Brain and Language*.
- Parisi D., Denaro D., & Cangelosi, A. (1997). Categories and word meanings are not single entities in the mind. *Technical Report NSAL-97010*, Institute of Psychology, CNR Rome.
- Saunders, G.M., & Pollack, J.B. (1996). The evolution of communication schemes over continuous channels. *Proceedings of the SAB'96 Conference on the Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press.
- Saunders, G.M., Angeline, P.J., & Pollack, J.B. (1994). Structural and behavioral evolution of recurrent networks. In D.J. Cowan, G. Tesauro, & J. Alspector (eds.), *Advances in Neural Information Processing 6*, Morgan-Kaufmann.
- Sharkey, N.E., & Jackson, S.A. (1994). An internal report for connectionists. In R. Sun (ed.), *Computational architectures integrating neural and symbolic processes*. Boston: Kluwer Academic Press.
- Steels, L. (1996). Self-organizing vocabularies. *Proceedings of the V Artificial Life Conference*. Tokyo, Japan
- Steels, L. (1997). Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation. In J. Hurford & Studdert-Kennedy (eds.) *Evolution of Human Language*. Edinburgh University Press.

- Steels, L., & Vogt, P. (1997). Grounding adaptive language games in robotic agents. In P. Husbands & I. Harvey (eds.), *Proceedings of the Fourth European Conference on Artificial Life*, Cambridge, MA: MIT Press.
- Webb, B. (1994). Robotic experiments in cricket phototaxis. *Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, MIT Press.
- Werner, G.M., & Dyer, M.G. (1991). Evolution of communication in artificial organisms. In C.G. Langton, C. Taylor, J.D. Farmer, & S. Rasmussen (eds.) *Artificial Life II*. Reading, Mass., Addison-Wesley.
- Werner, G.M., & Dyer, M.G. (1994). BioLand: a massively parallel environment for evolving distributed forms of intelligent behavior. In H. Kitano (ed.) *Massively parallel artificial intelligence*. Cambridge, MA: MIT Press.
- Werner, G.M., & Todd, P.M. (1997). Too many love songs: Sexual selection and the evolution of communication. In P. Husbands & I. Harvey (eds.), *Proceedings of the Fourth European Conference on Artificial Life*, Cambridge, MA: MIT Press.

Figure Captions

Figure 1. Neural networks controlling the behavior of speaking and listening organisms.

Figure 2. Average fitness across 1000 generations of three different populations: without language, with externally imposed language, and with evolved language. Each curve is the average of 5 different replications of the same simulation.

Figure 3. Typical frequency distribution of the 8 possible signals produced by all the individuals in each of 10 generations (generation 0, 100, 200, etc., up to generation 1000) in one replication of the simulation with the population with evolved language.

Figure 4. Evolutionary increase in average fitness in a population with no language. Also shown in the figure is the evolutionary change in an index of the quality of the language inadvertently produced by the population.

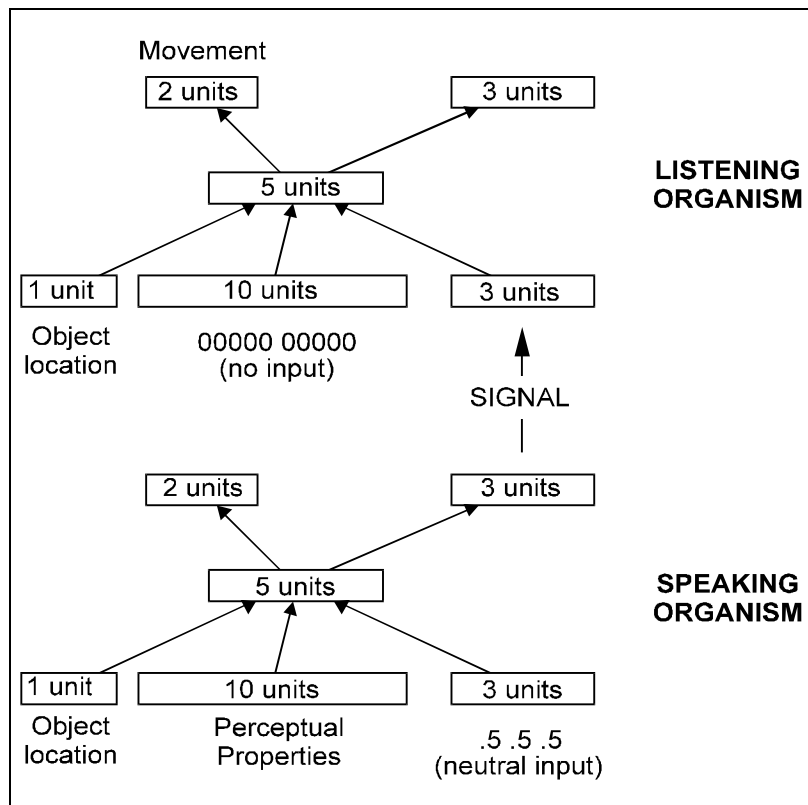


Figure 1




Figure 2





Figure 3




Figure 4