

An Artificial Life Approach to Language

Domenico Parisi

Institute of Psychology, National Research Council, Rome, Italy

The aim of the paper is to show that an Artificial Life approach to language tends to change the research agenda on language which has been shared by both the symbolic paradigm and classical connectionism. Artificial Life Neural Networks (ALNNs) are different from classical connectionist networks because they interact with an independent physical environment; are subject to evolutionary, developmental, and cultural change, and not only to learning; and are part of organisms that have a physical body, have a life (are born, develop, and die), and are members of genetic and, sometimes, cultural populations. Using ALNNs to study language shifts the emphasis from research on linguistic forms and laboratory-like tasks to the investigation of the emergence and transmission of language, the use of language, its role in cognition, and language as a populational rather than as an individual phenomenon. © 1997 Academic Press

1. INTRODUCTION

Language has been one of the success stories of the symbol manipulation paradigm in cognitive science. Chomskian linguistics, “box-and-arrow” psycholinguistics, and neurolinguistics have discovered and clarified many previously little-known aspects of language (e.g., the regularities underlying complex syntactic structures, how the mental lexicon is organized, and new types of linguistic disturbances) and have introduced theoretical ideas and modeling tools successfully applied not only to language but also to many other cognitive phenomena as well. It is not surprising, therefore, that language is seen as a great challenge for the connectionist paradigm which, in many respects, aims at replacing the symbol manipulation paradigm by proposing brain-inspired theoretical models (neural networks) and quantities rather than symbols as the basic entities manipulated in such models. In response to this challenge, connectionists have attempted to deal with linguistic phenomena and to propose alternative accounts of these phenomena since the beginning of connectionism in the 1980s (McClelland & Rumelhart, 1986; Sejnowski & Rosenberg, 1987; Seidenberg & McClelland, 1989;

Address reprint requests to Domenico Parisi, Institute of Psychology, CNR 15, Viale Marx, 00137 Rome, Italy. Fax: 39-6-824737. E-mail: domenico@kant.irmkant.rm.cnr.it.

for more recent connectionist treatments of language cf. Elman, 1990, 1993; Reilly & Sharkey, 1992), and various connectionist models of language disorders have been published (Hinton & Shallice, 1991; Patterson, Seidenberg, & McClelland, 1989).

Connectionist simulations of language and language disorders have yielded interesting results that challenge traditional interpretations of important linguistic phenomena. For example, it has been shown that systematic linguistic behavior can be accounted for without invoking the possession and use of rules (Rumelhart & McClelland, 1986; Plunkett & Marchmann, 1991) and that linguistic categories can be extracted from linguistic experience without postulating an innate basis for these categories (Elman, 1990). In the field of language disorders it has been shown that connectionist models can account for seemingly puzzling phenomena such as the co-occurrence of semantic and visual errors in deep dyslexia (Hinton & Shallice, 1991; Plaut & Shallice, 1993) and can suggest alternative models of reading (e.g., single-route vs. double-route models) based on connectionist simulations of reading disorders (Patterson, Seidenberg, & McClelland, 1989). (For a review of connectionist simulations on language, cf. Elman, 1995.)

The aim of the present paper, however, is not to judge the validity of these connectionist claims or to enter into the continuing debate on whether connectionism can answer the challenge represented by complex syntactic phenomena and by linguistic systematicity more generally (Fodor & Pylyshyn, 1988; Hadley, 1994). Rather, we will argue that the connectionist critique of symbol manipulation accounts of language has not been radical enough. Connectionism appears to have inherited the research agenda of the symbol manipulation paradigm with respect to language in that it tries to give alternative answers to the same research questions posed by the symbolic paradigm. We think, on the contrary, that the research agenda should be changed and that connectionism has the potential to pose and answer important questions concerning language that have been neglected during the period of prevalence of the symbol manipulation paradigm. However, in order to express this potentiality of connectionism, what we will call "classical" connectionism must be partially revised and connectionism should be viewed in a new perspective, that of Artificial Life. In this paper we will briefly outline an Artificial Life approach to connectionism and we will try to show how this approach can be used to address different aspects of language than those studied by both classical linguistics/psycholinguistics and classical connectionist simulations so far.

2. AN ARTIFICIAL LIFE APPROACH TO CONNECTIONISM

Neural networks are the theoretical and simulation tools used by connectionists to model behavior and cognition. Neural networks are inspired by the nervous system. Network units resemble neurons. Connections between

units resemble synaptic junctions between neurons and the weights on the connections resemble the strength of these junctions. The activation level of units is similar to the firing rate of neurons and the manner in which the activation level of a unit is determined is similar to the integrative function of the nervous system (the sum of excitations and inhibitions arriving to the unit). Finally, the strength of connections can be changed and connections can be created or eliminated reflecting the plasticity of the nervous system (learning).

Artificial Life (Langton, 1992) is defined as the attempt to study all biological phenomena by reproducing them in artificial systems (e.g., by simulating them in a computer). Since nervous systems are biological entities, neural networks by definition fall within the classification of Artificial Life. However, if we explicitly look at neural networks in the perspective of Artificial Life, simulations using neural networks take a different form with respect to classical connectionist simulations. We will consider three aspects in which neural networks viewed in an Artificial Life perspective (Artificial Life Neural Networks, or ALNNs) differ from networks used in classical connectionist simulations. These three aspects are (a) the role of the environment, (b) the different types of change that occur in organisms, and (c) the view of ALNNs as (part of) whole organisms and populations of organisms.

Classical connectionism has an impoverished notion of the environment. The environment tends to be equated with the input patterns that the researcher imposes to the network's input units at any given cycle (cf., e.g., McClelland, 1989, p. 14), but the environment has no independent physical existence. (An exception to this is the use of neural networks to control dynamical systems, c.f. Miller, Sutton & Werbos, 1992.) Biological entities of all kinds, on the other hand, constantly exchange energy and information with the external environment. Therefore, Artificial Life naturally tends to have a richer and more realistic notion of environment. The environment exists physically outside the organism and the input patterns arriving at the network's input units are an effect of this independent environment. It is the environment, not the researcher, that decides what the next input to the network is going to be. Furthermore, the organism (the network) acts on the environment and, therefore, by changing the environment or its relation to the environment, it at least partly determines its own future input. In conclusion, while classical connectionist simulations are only concerned with what happens inside a neural network, what is simulated in ALNN simulations is what happens both inside the network and inside the environment and, especially, in the interactions between network and environment.

Change is a crucial property of all biological phenomena. Organisms are subject to a variety of types of change and understanding organisms (more specifically, their phenotypic trait called behavior) requires understanding all types of change and how each type interacts with the others. Classical connectionism studies only one type of change, i.e., learning. Learning is

change occurring in individual organisms during their life which is mainly due to the particular experience of the individual in the particular environment. Other forms of change are biological evolution, biological development, and cultural evolution. Biological evolution occurs in populations of organisms and is due to the selective reproduction of the genotypes of the best individuals and the constant addition of variability to the gene pool through sexual recombination and random changes in the inherited genotypes (mutations). Biological development is change in the individual organism that is mainly due to the execution of "instructions" encoded in the individual's genetic material. Cultural evolution occurs only in some populations of organisms, e.g., humans, and it concerns information transmitted not genetically but somatically through learning from others or from artifacts made by others.

While classical connectionism is mostly restricted to learning, simulations using ALNNs tend to address all forms of change affecting organisms and, especially, the interactions among the different forms of change. Evolutionary change is studied within AL using genetic algorithms (Holland, 1975; Goldberg, 1989). Developmental change is addressed by modeling the temporal process of mapping the inherited genotype into the phenotypic neural network (cf., e.g., Nolfi & Parisi, 1995). Cultural change can be modeled by having neural networks learn by imitating the behavior of other networks and examining the changes that occur in the transmission of behavior from one generation to the next (Hutchins & Hazelhurst, 1995; Denaro & Parisi, 1995).

Finally, while classical neural networks tend to be viewed as abstract input/output systems, ALNNs are models of the nervous system of organisms that have a body (with an external morphology, internal organs beyond the nervous system, a genotype controlling the organism's development, etc.), have a life (are born, develop, and die), and are members of genetic and, in some cases cultural, populations made up of successive generations.

Studying language using ALNNs rather than classical neural networks has at least three consequences. Classical connectionist simulations tend to simulate laboratory-like tasks and situations while simulations with ALNNs are more likely to address the real (although hugely simplified) ecologies in which language has evolved in the population and is acquired and used by the individual. Classical simulations tend to study the learning of linguistic forms while simulations with ALNNs tend to combine the linguistic forms with their referents and meanings and to study the cognitive and pragmatic components of language. Language is viewed as a mental phenomenon taking place inside an individual in classical connectionism while it is studied also as a social and cultural phenomenon in ALNN simulations.

In the remaining sections of this paper we will describe some connectionist simulations of language phenomena in which connectionism is viewed in an Artificial Life perspective. It will be abundantly clear that only some limited

aspects of language have been addressed so far in this new perspective. Hence, in some cases we will be forced to describe possible (imagined) simulations rather than simulations that have already been done. However, even the small amount of work that has been done shows the promise of ALNNs for changing the research agenda for the study of language.

3. SIMPLE ANIMAL COMMUNICATION

Imagine a population of organisms living in a shared environment. The organisms exhibit various type of behaviors through which they interact with each other and with the other entities present in the environment. Among the behaviors exhibited by the organisms are communicative behaviors. Communicative behaviors are behaviors that are exhibited with the goal to generate particular sensory inputs for other organisms, mostly conspecifics. "Goal" is defined here in purely evolutionary terms. An individual has the goal to generate a particular behavior if the individual is part of an evolving population in which the individuals that tend to exhibit that behavior have more reproductive chances than individuals that do not exhibit the behavior. The behavior in our case is some motor action that causes some sensory input (a signal) for another individual. The other individual responds to this input by producing some behavior on its part. This behavior also represents a goal for the responding individual in that individuals exhibiting this particular response to the signal are more likely to reproduce than individuals that respond in other ways or that do not respond at all. In other words, communicative behavior includes both producing signals and responding to signals. The two behaviors must coevolve. The behavior of producing signals in the appropriate circumstances increases the reproductive chances of the producer only if the other individual responds appropriately to the signal at the same time.

Werner and Dyer (1991, 1994) have simulated the evolution of communicative behaviors in populations of neural networks. One (female) individual exhibits some behavior that causes some sensory input for another (male) individual. The female generates the signal for the male in response to its perception of the male. The male cannot directly perceive the female but it responds to the signal by being guided by the signal to approach the female. When the two individuals come together they can mate and have offspring. The initial population includes neural networks with random connection weights. Hence, both female and male individuals that are members of the initial population do not respond to sensory input (perception of the male for the female and perception of the signal for the male) with the appropriate output (generation of the signal for the female and approaching the source of the signal for the male). However, selective reproduction and the constant addition of variability through mutations and sexual recombination cause the progressive emergence in the population of individuals exhibiting the appropriate communicative behaviors.

When some particular communicative behavior tends to be regularly exhibited by the producer in the presence of some particular object (or action or event or property), the object becomes the “referent” of the signal caused by the behavior. The object is perceived by the input units of the producer’s neural network when the communicative behavior is produced and is perceived by the input units of the receiver’s neural network together with the signal. As a consequence, the neural representation (say, the activation pattern of some hidden network units) of the object tends to be associated/integrated with the neural representation underlying the signal-producing behavior in the producer’s neural network and with the neural representation of the sensory signal in the receiver’s neural network. These associated/integrated internal representations become the “meaning” of the signal. (We use the awkward expression “associated/integrated” because we want to leave it open if the two neural representations remain separate and they simply tend to evoke each other through their interconnections, or the two representations become integrated in a single distributed representation.) Notice that since the producer may also be a receiver of its own signals, e.g., in the case of auditory signals, and, furthermore, one and the same individual can act as both producer and receiver of the same signal in different occasions, the cluster of associated/integrated neural representations in the neural networks of all individuals tends to include the neural representation of (a) the signal as sensory input, (b) the referent of the signal, and (c) the signal-producing behavior.

In more recent simulations Werner (unpublished paper) has shown how neural networks can learn to generate the same behavior in response to both an object, and a signal that has the object as its referent. Using a reinforcement learning procedure, a neural network can learn to respond, for example, to a signal “hawk” emitted by a conspecific by running away from the source of the signal as it would run away from an actually experienced hawk. Furthermore, a behavior that is generated in response to some signal can be transferred to another signal if the two signals are experienced together. Using a Hebbian learning procedure, the same network can learn to respond by running away from a snake if it perceives the snake together with the signal “run-snake” while it did not run away from the snake prior to learning.

4. LANGUAGE AS AN AID IN CATEGORIZATION

One important function of human language results from the fact that organisms are often required to respond in the same way to different objects (or views of objects), actions, events, etc. In other words, they must generalize across different inputs by generating the same output in response to all inputs belonging to the same category. In order to do so they must recognize a particular input as belonging to the category. Now, imagine that the same




FIG. 1. A population of organisms living in an environment that contains food elements (round objects) and dangerous elements (angular objects).

signal is paired with different referents in the experience of both producers and receivers. If an organism must respond in the same way to different objects, actions, etc., experiencing these different objects, actions, etc., in association with an identical signal, can make it easier to recognize the different objects, actions, etc., as requiring the same response. In these circumstances, of course, the meaning of the signal, i.e., the neural representation of the referent of the signal, may become a complex entity reflecting (a) particular features of the different referents paired with the same signal (meaning as family resemblance); (b) a prototype referent never encountered in reality (meaning as prototype); (c) a set of abstract features shared by the different referents (meaning as abstraction).

Imagine an organism living in an environment that contains two types of objects: foods and dangers (cf. Fig. 1). Both foods and dangers are defined in evolutionary terms. Foods are those objects such that when the organism happens to step on one of them the individual's reproductive chances are increased. Dangers, on the contrary, are objects that decrease the reproductive chances of the individual if the individual happens to step on them. Food objects are not all identical objects, and the same is true for dangers. Hence, the individual must possess two different abilities to survive and reproduce




FIG. 2. Performance of the best individual of the last generation tested with old patterns, new patterns, and prototypes in the simulation with labels and in the simulation without labels.

in this environment. It must be able to recognize an object as either food or danger and it must be able to associate the appropriate response to each category of objects, i.e., it must approach foods and avoid dangers.

Neural networks can be shown to be able to acquire these two abilities (Floreano, Miglino, & Parisi, 1991). Food objects are encoded as binary patterns deviating in one or two bits from a prototype pattern. Danger objects are encoded as deviations from a different prototype pattern. At any given time a neural network is informed by its senses about the perceptual properties (binary pattern) of the nearest object, either food or danger, and it responds by turning or moving in the environment. Networks belonging to the later generations in an evolving population of networks tend to approach foods and to avoid dangers.

However, the evolutionary emergence of these abilities turns out to be rather slow and difficult. Now imagine that when a network happens to perceive a food object the network also receives some additional sensory input (the signal "food") which is identical for all food objects. Similarly, when the network perceives a danger it also receives some additional sensory input (the signal "danger") which is identical for all dangers but is different from the signal accompanying food objects. (The two signals are encoded as the binary patterns 01 for food and 10 for danger in two additional input units.) It can be shown that the behavior of approaching foods and avoiding dangers reaches higher performance levels in these circumstances. Figure 2 shows the fitness (the number of food objects minus the number of danger objects the individual happens to step on during its life) of individuals tested with the patterns used during evolution (old patterns), new patterns not used during evolution (new patterns), and the two prototype patterns for food and danger

themselves (prototype patterns) in the simulation with linguistic signals and in the simulation without signals.

Language, then, can function as an aid to categorization by stabilizing the class of objects that must be put in the same category and responded to in the same way. As we will suggest in the next section, this is likely to happen not only when the linguistic signal comes from outside—as in the simulations just described—but also when it is self-generated by the same individual who must assign an input to a category and respond to the input based on that category.

5. CULTURAL TRANSMISSION OF LANGUAGE

In the food and danger simulations the organism acts only as a receiver of signals and it evolves an ability to respond appropriately to these signals. It is interesting to ask, however, where these signals come from. One can suppose that there is another individual in the local environment of the first individual and that this second individual generates the appropriate signal (“food” or “danger”) when it perceives either a food object or a danger object. In other words, both the first and the second individual perceive the same particular food or danger object. The second individual reacts by generating the signal “food” or “danger” as is appropriate and the first individual perceives the signal generated by the second individual together with the food or danger.

Why should the second individual bother to generate signals in the presence of the first individual? The evolutionary “goal” of the first individual is quite clear. Individuals who respond to the signal “food” (“danger”) by approaching (avoiding) the object they currently perceive are more likely to reproduce than individuals who do not do so. Hence, the evolutionary emergence of an ability to understand these signals, i.e., to respond appropriately to them, can be explained by our simulations. But why should individuals who perceive food or danger objects in the presence of another individual develop a tendency to respond by emitting the signal “food” or “danger”? In Werner and Dyer’s simulations described above, females generate signals for males because this behavior increases their chances of mating with males. But what is the adaptive advantage, for the generator of the signal, of generating signals benefiting the other individual in the case of our food and danger simulations?

The answer may be found in kin selection theory. According to kin selection theory (Hamilton, 1964) individuals tend to evolve altruistic behaviors that benefit another individual at a cost to the individual emitting the behavior if the two individuals are genetically related and the benefit is higher than the cost. By emitting the altruistic behavior the individual reduces the reproductive chances of its own genes but it increases the reproductive chances of its relative’s genes. The relative’s genes are similar to the altruist’s genes and, therefore, they are likely to incorporate the same altruistic tendencies.

In other simulations (Parisi, Cecconi, & Cerini, 1995) it has been shown that neural networks can evolve altruistic behaviors toward their kin. The networks live in an environment that contains food elements and their reproductive chances (fitness) depend on the number of food elements they are able to eat. Their life history is divided up into two successive stages: reproductive immaturity and adulthood. During immaturity individuals are unable to find the food they need to survive and depend on the food given to them by their parents. Individual networks inherit an "altruism gene" that can take a value from 0 (total egoism: parents eat all the food they find) to 1 (total altruism: parents give all the food they find to their offspring). The population evolves an appropriate intermediate value for the altruism gene which allows parents to eat enough food to survive and generate more offspring and to give enough food to their offspring to allow them to survive and reach maturity. By giving food to their offspring parents decrease their individual reproductive chances but increase the reproductive chances of individuals (their offspring) having their same altruistic tendencies (same value of the altruism gene).

If we assume, then, that the cost to the individual emitting the food signal in the presence of both a food object and one of its genetic relatives is lower than the benefit to the relative being provided with a signal that helps it recognize the object as food, we can expect the signal-emitting behavior to evolve in the population. In real organisms (humans), these circumstances may occur when parents name the objects (actions, etc.) they perceive in the presence of their young offspring. They are helping their offspring to categorize reality in useful ways and the evolutionary advantage for the offspring is greater than the cost for the parent.

Of course, children can also be helped in their categorization of entities in the world by the language produced by nonrelatives, because language is produced by both relatives and nonrelatives for their own goals, i.e., to increase their own individual fitness, and the child may be helped in its categorization efforts by noticing what entities in the world are associated with (categorized by) which signals produced by other individuals that already know the language. What is more critical is that children cannot always count on others to provide them with useful signals for categorizing reality. They may be confronted with reality when they are all alone. Who is going to provide them with linguistic signals to help them categorize reality in these circumstances?

Our simulations have an even more crucial limitation. We have assumed so far that both the behavior of producing signals and the behavior of responding to signals are biologically evolved behaviors. In our simulations individual networks are born with genetically inherited tendencies (incorporated, e.g., in their inherited connection weights) to produce certain signals in certain circumstances and to respond to each signal with the appropriate behavior. But (human) linguistic behavior is not genetically inherited. It is

learned—although probably on the basis of various genetically inherited constraints and predispositions.

In fact, language is not only learned but it is also learned from others and is culturally transmitted from one generation to the next. Hence, if we want to pursue our research program applying an Artificial Life perspective to the study of language, we must be able to incorporate learning from others and cultural transmission into our simulations.

A simple neural network model of learning from others, based on a model proposed by Hutchins and Hazelhurst (1995), is the following. Imagine two networks. One network already knows how to respond to a set of inputs; this is the “model.” The other network does not know how to respond to those inputs; this is the “cultural learner” or, more specifically, the “imitator.” The two networks are presented with the same input. Both networks compute their respective outputs. Then, the imitator uses the output of the model as its teaching input for learning using the backpropagation procedure (Rumelhart, Hinton, & Williams, 1986). The imitator changes its connection weights in such a way that the discrepancy between its output and the model’s output is progressively reduced. After a certain number of learning cycles, the behavior of the imitator will be similar to the behavior of the model. The imitator will have learned whatever ability (responding to input with the appropriate output) was already possessed by the model.

We can apply this model to our situation in which an adult network generates a particular signal when both a particular object (the signal’s referent) and one of its offspring are simultaneously present. To change this situation into a language learning situation we must modify the network architecture of the offspring by adding some new output units similar to the signal-producing output units of the adult network. Both the parent and the offspring generate a signal in response to the object, but at the beginning of learning the offspring’s signal is likely to be different from the parent’s signal. The offspring uses the parent’s signal as teaching input to change its connection weights in such a way that in future occasions its signals in response to particular objects will be closer to the parent’s signals. After a certain number of learning cycles, the offspring has learned the repertoire of signals (the language) of its parent. The language has been transmitted from parent to offspring, from one generation to the next.

Another function of generating signals in the presence of both an object and another individual is to allow the other individual to learn to self-generate the signals. If kin selection theory can explain why parents generate signals in the presence of their offspring because this helps the offspring categorize reality, another function of the behavior of the parent appears to be to allow parents to transmit the repertoire of signals to their offspring—which guarantees the survival of their genes on a longer time perspective.

But learning the language from one’s parents has two useful consequences, not only one. An individual that has learned early in life to generate the

appropriate signals in the presence of objects will be able, as we have already seen, to help its offspring both to categorize reality in useful ways and to learn these signals in the same way as the individual has learned them from *its* parents. But, in addition, learning early in life to generate the appropriate signals for objects can help the same individual that generates the signals to categorize reality even in the absence of other individuals providing the signals. The results of our simulations (cf. Fig. 2) have shown that signals can help individuals recognize objects and respond to these objects appropriately when the signals are provided to them from outside (by another individual). However, it might also be advantageous to self-generate signals in response to objects in order to have these self-generated signals mediate the individual's own behavior with respect to the objects. (We intentionally use the term "mediate" here to refer to the mediational role of language in behavior stressed by Vygotsky and other psychologists of his school; cf. Diaz & Berk, 1992). If a perceived object must be recognized as food in order to respond to the object by approaching it, it might be useful to first name the object by self-generating the signal "food" and then respond to both the object and the self-generated signal. The signal can help recognize the object and respond appropriately.

One might object that if an individual must acquire the ability to respond with the same behavior to different inputs, it is not clear why generating the same approaching behavior in response to different food objects is more difficult than generating the same food signal in response to different food objects. What is the advantage of interposing a label between the food and the approaching behavior? One hypothesis that could be tested with simulations is the following. While the individual who has no signals must do a single mapping from foods (input) to the approaching behavior (output), the individual who interposes a signal between the food and the approaching behavior must do two mappings, one from foods to the signal and the second one from the signal to the approaching behavior. However, the direct "sensed food/approaching behavior" mapping required of organisms that cannot self-generate signals is many to many. There are many food objects and there are many different behaviors that are all approaching behaviors. On the other hand, for organisms that can self-generate signals the "sensed food/signal" mapping is many to one and the "signal/approaching behavior" mapping is one to many. (Obviously, the same signal tends to be realized in different ways in different occasions. But the dimensionality of the space of signals appears to be much more restricted than the dimensionality of the space of objects and behaviors.) Now, it is possible that for a neural network substituting two mappings, one many-to-one and the other one-to-many, in place of a single many-to-many mapping is a smart move that can explain why language for oneself has emerged and why language is useful for thinking. In fact, it is not clear if human language has emerged because of its utility for social communication and then it has been used as a preadap-

tation for talking to oneself, or vice versa. It is also possible, of course, that both evolutionary pressures have been present from the beginning. (Notice that language for oneself is useful not only for categorizing reality but also for many other cognitive operations; cf. Section 8 below and Diaz & Berk, 1992.)

6. HOW CAN LANGUAGE ORIGINATE IN A POPULATION WITH NO PRIOR LANGUAGE?

In a population of neural networks in which each new generation of individuals learns to label the objects of the environment by learning to imitate the labeling behavior of older individuals that already know the labels, the repertoire of labels or signals is culturally transmitted. Individuals are not born with the ability to generate the appropriate signals in the presence of various objects, nor does this ability simply develop as a result of maturation with a limited and aspecific role of the environment. The signals that are learned depend on the particular signals that are produced by the conspecifics in the environment of the individual. What may be genetically inherited by humans is a generic tendency to learn by taking the output of another network as one's own teaching input (learning by imitation), some time constraints on when the learning of language must take place (cf. Elman's simulations described in Section 8 below), and perhaps some other predispositions and constraints that are language-universal. However, the lexical signals of the specific language must be learned from others. The cultural, not genetic transmission of the repertoire of lexical signals is typical of human languages, although some role of learning by imitation in the acquisition of communicative behaviors can be found in other animals (e.g., in birds; cf. Munding, 1980).

It remains an open question, however, how the repertoire of signals is initially introduced in a population. Hutchins and Hazelhurst (1995) have done some simulations that shed some light on this issue. They show how signals for objects (actually, phases of the moon) can emerge in a population of neural networks that initially lack them in such a way that the signals for different objects are maximally different and signals for the same objects are maximally similar across the individuals comprising the population. The networks learn to generate signals for the different objects by using the signals generated by other randomly selected individuals as their teaching input (cf. the preceding section). No individual already knows the appropriate signals at the beginning of the simulation and in fact there are no intrinsically appropriate signals (i.e., signals are conventional). The repertoire of signals emerges spontaneously from the interactions (reciprocal learning) among the individuals of the population.

The networks in Hutchins and Hazelhurst's simulations must learn to generate an appropriate functional behavior in response to each phase of the

moon. (One such functional behavior might be, for example, to travel to the seaside when the moon is full because this is good for fishing.) The functional behavior appropriate for each phase of the moon is taught to the networks using the backpropagation procedure. (Actually, the networks are autoassociators that must generate an output identical to their input.) Hutchins and Hazelhurst interpose an intermediate “verbal” layer between the input and the output layers of their networks and this verbal layer is taught by using the pattern of activation of the verbal layer (verbal output) of other networks. As a result, the learning of the appropriate functional response to each phase of the moon (i.e., the autoassociation task) is facilitated, and a lexicon emerges in the population as a set of maximally shared and maximally distinctive verbal labels.

Hutchins and Hazelhurst describe various other suggestive results. For example, they show that consensus in the choice of signals among the individuals comprising the population is more difficult to achieve in larger than in smaller populations. On the contrary, distinctiveness of signals for different objects is easier to obtain with increasing population size. Furthermore, they have experimented with the introduction of new individuals (neural networks with random weights) in a population that has already reached a stable, shared lexicon. If the population size is large a newcomer will simply learn the shared lexicon. With smaller population sizes the arrival of a newcomer may destroy the lexicon previously shared by the population.

Hutchins and Hazelhurst suggest that the disruptive role of the arrival of new members (newborns) to a linguistic community may be blocked by the emergence of a critical period for language learning. Adults would be less and less prone to change their linguistic habits and this would shelter them from the risks of absorbing the disorganized linguistic behavior of newcomers. This critical period could be implemented in neural networks by progressively reducing the learning rate of growing individuals. (The learning rate of a neural network is a constant that determines how much the network’s connection weights are changed as a function of their contribution to the global error of the network, i.e., to the observed discrepancy between the output computed by the network and the teaching input.) During the critical period for language learning, young individuals would have a large learning rate that would allow them to absorb the language of the community. But as the individuals become older, the learning rate is decreased and there is less of a tendency (and capacity) to learn new linguistic habits.

Hutchins and Hazelhurst prefer not to implement the progressive reduction in learning rate as an arbitrary decision of the researcher, but they suggest that such reduction may be an inherited tendency resulting from biological evolution. In fact, one might imagine a population of neural networks in which different inherited genotypes specify different possible schedules of change in learning rate during the life of individuals. If the possession and maintainance of a shared and distinctive lexicon is advantageous from the point of view of fitness, evolved individuals of later generations are likely

to inherit genotypes that prescribe a progressive reduction of the learning rate during the lifetime of an individual, i.e., that implement a critical period for language learning blocking the disruptive action of newcomers to the language community.

Hutchins and Hazelhurst's model for the emergence of language (lexicon) ensures that the same language is shared by all individuals comprising the population since all individuals learn from all individuals. On the contrary, in the previous model in which offspring learn the language from their parents, a different language would be transmitted along each genetic line. However, one last result reported by Hutchins and Hazelhurst concerns the emergence of dialects. If the reciprocal learning of signals is panmictic, i.e., all individuals in the population are equally likely to be chosen as models for learning in any particular occasion, the population tends to remain linguistically homogeneous. If, on the other hand, models for learning are chosen selectively, for example on the basis of similarity in linguistic behavior (parallel to assortative mating in sexual reproduction), there is a tendency for the population to become linguistically heterogeneous and to include separate subpopulations speaking different dialects, i.e., different signals for the same objects in the environment.

Hutchins and Hazelhurst's model for the cultural emergence of language in a population of neural networks initially lacking it assumes that individuals find it useful to develop a shared lexicon because this helps them to respond appropriately to the input in an already well-defined task (the autoassociation task). (In fact, Hutchins and Hazelhurst's networks find it useful to interpose a linguistic label between the input and the output of their task; cf. the previous section). However, the cultural transmission of behaviors via imitation from one generation to the next can create previously inexistent capacities by itself, even if there is no previously defined task that the culturally emerging capacities can help accomplish. For example, an initial population of neural networks with randomly assigned connection weights can culturally evolve from zero a capacity to approach food elements present in the environment if the behavior is culturally transmitted selectively (i.e., only the best individuals act as models for the next generation) and some random noise is added to the teaching input provided by the model to the imitator (this random noise is analogous to mutations in biological evolution) (Denaro & Parisi, 1995). This model of cultural evolution could be applied to language by having a population of neural networks evolve a genetically inherited ability to approach food and avoid dangers (cf. Section 4) on one hand and a culturally inherited ability to generate the appropriate signals that can help in categorizing inputs as either food or danger on the other.

7. CULTURAL (HISTORICAL) CHANGE IN LANGUAGE

The same model of cultural learning discussed in the preceding section, i.e., learning from others by using their behavior (output) as a model to be

imitated, can be used to study historical linguistic change—a form of cultural change. Using the backpropagation procedure, Hare and Elman (1993) have taught a neural network to generate the past tense of English verbs on being given the present tense as input (Rumelhart & McClelland, 1986). As is well known, English verbs can be classified in different classes as a function of the particular nature of the present/past tense mapping (regular (*work/worked*), irregular (*teach/taught*), identical (*put/put*), etc.) with different classes having different type and token frequencies. Many simulations have been done using this paradigm yielding many interesting results and analyses (Rumelhart & McClelland, 1986; Plunkett & Marchmann, 1991, 1993) and in particular showing how type of mapping and both type and token frequency of verbs influence learning.

Hare and Elman's simulations begin with a network that is taught the present/past tense mappings of Old English verbs. The training is stopped before the network reaches perfect learning and the output of the network is used as teaching input to train a second network. This procedure is repeated for a number of networks in a sequence simulating successive generations of people learning English. The error of each network trickles down and accumulates in successive networks due to the imperfect learning in each network. Given the complex structure of the present/past tense mapping of English verbs, Hare and Elman find that the behavior of successive networks is subject to reorganization reflecting the historical changes that have actually occurred in the English verb system over time.

8. OTHER ASPECTS OF THE ONTOGENY OF LANGUAGE

Let us come back to how language is acquired by the individual. We have assumed that linguistic signals are learned by imitating the signals produced by other individuals who already know the language, and we have proposed a simple neural network model of learning by imitation. The learner uses the output of another network (the model) as its teaching input and learns to generate the same signals as the other network in response to the same referents.

This type of learning is more complex than it appears. It has two components: (a) learning to produce the *same* signals and (b) learning to produce these signals *in response to* the same referents. Children may not learn the two components at the same time. Unlike the sounds produced in the very first months of life, the sounds produced by a child after a few months tend to be similar to the particular sounds that can be heard in the child's environment. However, the child appears to learn to imitate the sounds produced by other individuals prior to learning to produce these signals in response to particular referents—which only happens toward the end of the first year. Furthermore, the task of learning to produce the same signals that are produced by the adults who already know the language is not a simple task for

the following reason. Our model of learning by imitation assumes that the learner network (child) compares its own output with the output generated by the model network (adult) in order to discover how the two outputs diverge and to use the divergence to modify its connection weights. However, while the output of the child network encodes phono-articulatory movements, the output of the adult network is not directly accessible to the child as phono-articulatory movements but only as the sounds that are physically caused by the adult's phono-articulatory movements. Hence, it is not clear both technically and conceptually how the child network can compare two entities that are described using different dimensions (e.g., degree of opening of the mouth for phono-articulatory movements and type of formants for sounds).

This is a general problem for learning to imitate. If learning to imitate consists in comparing one's output with the teaching input resulting from the output of a model in order to modify one's weights to gradually approximate the output of the model, the learner's output and the teaching input must encode the same type of information. But while in classical connectionism one could *assume* they encode the same information (as in the case of Hutchins and Hazelhurst's simulations), ALNNs are constrained by what input or teaching input the environment actually makes available to the learner's network. And in learning to imitate the learner's output encodes motor acts while the teaching input encodes sensory information resulting from the model's motor acts.

To solve this problem for language learning, one possibility is to assume that the child first learns to map its own phono-articulatory movements into the sounds produced by these movements and then it can learn to generate phono-articulatory movements that map into sounds similar to the sounds that are produced by adults. (For a general account of the role of "forward models" in adaptive systems, see Jordan & Rumelhart, 1992). This can be modeled by first having a neural network learn to predict the sounds resulting from its own planned phono-articulatory movements. The network's input units encode some planned phono-articulatory movement and the network generates a description (prediction) of the sound that will be heard after the phono-articulatory movement has been physically executed (cf. Fig. 3). The phono-articulatory movement is then physically executed and the actual sound that is produced by this phono-articulatory movement is used by the network as teaching input to change its connection weights. After a certain number of learning cycles the network is able to predict correctly the sounds resulting from its own phono-articulatory movements. This is Stage 1 and it corresponds to the child generating various sounds in the first months of life (babbling) with the only purpose in learning is to predict the auditory consequences of its own phono-articulatory actions.

The child then moves to Stage 2 in which it learns to imitate sounds that are produced in the child's environment. The network architecture for this




FIG. 3. Network architecture for learning to predict the sounds resulting from planned phono-articulatory movements.

stage is shown in Fig. 4. The upper portion of the new architecture is the same network in Fig. 3, with the connection weights that have been found at the end of prediction learning in Stage 1. The lower portion is new and has randomly been assigned connection weights. The input units of the new architecture encode heard sounds (produced by adults in the child's environment). The lower portion of the network generates some phono-articulatory movements in response to these heard sounds, but before the phono-articulatory movements are physically executed the upper portion of the network




FIG. 4. Network architecture for learning to imitate heard sounds.




FIG. 5. Network architecture for learning to name objects.

generates a prediction concerning the sounds that will result from these phono-articulatory movements (using the prediction ability acquired in Stage 1). The predicted sounds are compared with the heard sounds and the discrepancy is used to modify the connection weights of the lower portion of the network. (The connection weights of the upper network are frozen in Stage 2.) After a number of learning cycles the network has learned to imitate heard sounds, i.e., to generate phono-articulatory movements that result in sounds similar to those heard by the child.

Stage 2 corresponds to the child learning to imitate the sounds produced by the adults in the child's environment in the second semester of life. In Stage 3, toward the end of the first year, the child begins to learn to associate sounds to referents. The network architecture for Stage 3 is shown in Fig. 5. A new piece of neural structure has been added to the network architecture. This new structure allows the network to encode the perceptual properties of referents and to learn to produce signals (sounds) in response to both heard signals (imitation) and referents (naming).

This three-stage model has been used to simulate the acquisition of the ability to imitate the nine vowel sounds of English and to use these nine vowel sounds to name each of nine different objects (Floreano & Parisi, 1994). The nine vowel sounds are described using the two main formants of each vowel while the phono-articulatory movements are described using Stevens and House's model of the phono-articulatory tract with three dimensions: (a) distance from glottis to point of maximum constriction; (b) radius

at point of maximum constriction; (c) opening of the mouth (Stevens & House, 1955).

The model of the early stages of language learning that we have described suggests ways in which the child may use its experience during its first year of life to first learn to predict the sounds resulting from its phono-articulatory movements, then to imitate the sounds produced by adults in the child's environment, and finally to associate these sounds with referents. We have assumed that this learning occurs in three consecutive stages to reflect the empirical evidence (although the learning of one stage can be continued in the next stage(s) both in reality and in our model). But our model predicts that a "babbling" stage (more or less random generation of motor actions with the purpose of learning to predict their sensory consequences) will precede imitation not only in language learning but also in other domains, and that the sounds produced by the child in Stage 1 will be different from the sounds produced in Stage 2.

The role of the child's experience in this learning is clear. But genetic-maturational factors are also likely to be crucial, for example, in causing the changes in network architecture from Stage 1 to Stage 2 to Stage 3 and the different types of learning that are made possible by these differently developing architectures. The evolution of network architectures that develop during the lifetime of the organism can be studied using the model described in Nolfi and Parisi (1995a) in which the inherited genotype specifies the parameters of axonal growth of a set of network units and the phenotype is a succession of developmental forms of the resulting neural network. (The model of the genotype-to-phenotype mapping is described in Nolfi & Parisi, 1995b and, in a more complex version that includes cell division and migration, in Cangelosi, Parisi, & Nolfi, 1994). Given this model, it might be possible to show that learning a language may be favored by the evolution of the particular sequence of developmental stages schematized in Figs. 3 to 5.

Another aspect of the interaction among evolution, development, and learning in the acquisition of language is suggested by some simulations done by Elman (1993). A neural network learns to predict the next word in a sentence by being given the current word as input. The network architecture includes some "context" units where a copy of the activation pattern of the network's hidden units is stored in each cycle. This activation pattern is input to the network's hidden units in the next cycle together with the current input (word). Hence, the network learns to predict the next word in a sentence using both the current word and a cumulative trace of the preceding context.

If the training corpus of sentences includes both syntactically simple (short) sentences and syntactically complex (long) sentences, the network has some difficulty learning the task. Syntactically complex sentences may require that the network takes into consideration long distance dependencies to predict the next word, and the network's contextual memory mechanism

(the context units) may be unable to do so. Elman has shown that the problem can be solved if the network in a first learning stage is given only simple (short) sentences and then in successive stages more complex (longer) sentences are added to the training corpus. However, this solution appears to be ecologically implausible, because the child appears to be exposed to all sorts of short and long sentences from the beginning. Another solution is to manipulate the contextual memory mechanism of the network. The network is exposed to the entire corpus of short and long sentences from the beginning, but its memory span changes during learning. In a first stage the activation pattern of the context units is erased each two or three cycles. Hence, in this stage the network can only use information about the very recent context (two or three preceding words) to predict the next word. In a second stage the memory span is increased by erasing the content of the context units each three or four cycles, then each four or five cycles, and so on for a number of successive stages. In these conditions, the network is able to learn to predict the next word in all sorts of sentences. The network (child) is exposed to all sorts of short and long sentences from the beginning of learning but it manages to solve the problem because its increasing memory span acts as a filter allowing the network to deal with short dependencies first and then with longer and longer dependencies as learning proceeds.

In Elman's simulations it is the researcher who manipulates the memory span of the neural networks by deciding at which cycle to erase the content of the context units. But Elman's networks can become members of an evolving population of ALNNs that inherit in their genes a developmental schedule for changing the memory span of individual networks during their life. In this case too, the reproductive advantages of being able to predict the next word in a heard or planned sentence may cause the evolutionary emergence of developmental schedules that gradually increase the memory span of the networks so that the networks can learn to make the correct predictions even if they are exposed to all sorts of sentences since the beginning of learning.

9. FROM SIMPLE SIGNALS TO LANGUAGE

We have occasionally used the word "language" to refer to the behavior of networks in the simulations described so far. But it is obvious that human language is much more complex than the simple signals produced or understood by those networks. In almost all simulations we have described signals are in one-to-one correspondence to the objects of the world, although this limitation is overcome in the simulations on labeling food and danger where signals are in a one-to-many correspondence with referents. But the human language is even more complicated than that. One and the same signal, e.g., "red," can be paired with two different categories of objects, e.g., red cars and red balls. We say that these are two different categories of objects because in other occasions red cars and red balls may be associated with two

different signals, respectively, “car” and “ball.” In fact, it is not only the case that the same signal can be paired with two different categories of objects but, conversely, two different signals can be paired with the same object, e.g., the signals “red” and “car” with a red car.

One direction of future research will explore if neural networks can learn to make these more complex associations without being confused by the training data. It is clear that it is because of these more complex signal/referent associations that language can play a role in fostering not only recognition and categorization (cf. Section 4 above) but also abstraction and selective attention. A network can be helped to abstract the perceptual property of being red from a perceived red car if the network is exposed to the signal red while perceiving the red car. Abstraction here means that the network is more ready to make its output dependent on that property of the car rather than on its other properties. Or, given a perceived scene that includes both a car and a ball, the network can selectively attend to, e.g., the car if the network is simultaneously exposed to the signal “car” rather than to the signal “ball.” Selective attention here means that the network is going to make its subsequent behavior (output) dependent on the portion of the input that encodes the car rather than on the portion encoding the ball.

Of course, organisms can recognize, categorize, abstract, and selectively attend to objects and properties even in the absence of language—and these abilities can be studied using classical neural networks. ALNNs are specifically interested in how language can enhance these cognitive abilities and how this role of language can explain its biological and cultural emergence.

An even more difficult test for neural networks will be the processing of complex signals (phrases or sentences) that are temporal sequences of simple signals (words). Ecological networks that produce temporal sequences of signals learn to produce each particular signal in response not only to external referents but also to the previous self-generated signals in the sequence. The model of imitation described in Section 5 can be extended to the learning of sequences of signals. The model assumes that the learner (imitator) is physically near to the model so that the learner can experience and respond to the same sensory input as the model. This shared sensory input will include both the external referent and the signal generated by the model in response to the external referent and to its own previous signal(s). At the beginning of learning the signals generated by the learner will be rather different from the signals generated by the model and, as a consequence, it is appropriate for the learner to produce each new signal in the sequence in response to the model’s previous signals rather than in response to its own preceding signals. However, as the ability of the learner to produce signals similar to the model’s signals increases, the learner can begin to make each successive signal depend on its own self-generated signals rather than on the signals generated by the model. In other words, sensory control of the learner’s linguistic output is progressively transferred from the model’s linguistic behav-

ior to the learner's own behavior. This is of course crucial if the learner must be able to produce sequences of signals in the absence of the model.

But producing and understanding sequences of signals involves much more than just learning to produce the same sequences of signals that are produced by other speakers. Complex signals have complex meanings and it remains to be seen if neural networks will be able to produce and understand sequences of signals in which the meaning of each signal is combined in systematic ways with the meaning of the other signals and there are special signals (function words and bound morphemes) that have the function to help in combining meanings together. It is possible, however, that this much discussed problem will become more tractable if it is dealt with using ALNNs rather than classical networks. Unlike classical neural networks ALNNs are intrinsically sensitive to temporal ordering since for ALNNs the input of each successive cycle in a sequence of input/output cycles necessarily includes the sensory consequences of the network's output in the previous cycle. Therefore, for ALNNs the sequence AB is necessarily different from the sequence BA. When element B is produced (or received) as part of the sequence AB, the input will include the sensory consequences of the network's production of A (or of the network's response to A)—which is not the case if element B is produced (or received) as part of the sequence BA. (This sensitivity to temporal ordering can be obtained in classical networks by adding some special units that record the network's state in the preceding cycle and re-input the state in the next cycle, as in Elman's simulations; cf. Elman, 1990.) Sensitivity to temporal ordering may be a critical factor in dealing with syntax.

Another consideration is that the generative character of language, i.e., the ability to produce and understand new combinations of previously known signals, may be more easily demonstrated in the behavior of ALNNs because unlike classical networks ALNNs acquire and use a language in an environment and, therefore, they can have referential (objects in the environment), cognitive (neural representations of these objects in the neural network), pragmatic (goals in producing and understanding signals), and social (what the other people do with language) support when they are asked to demonstrate their generalization abilities from old to new combinations of signals.

Aside from the ability of ALNNs to produce and generate complex linguistic signals, an Artificial Life approach to language is interested in studying how a complex language can emerge from a simpler language in a population of neural networks. It is possible that complex signals that are sequential combinations of simple signals arise when the repertory of signals that are needed by complex organisms with complex social interactions becomes too large to be acquired by learning each signal individually. This may represent an evolutionary pressure, at both the biological and the cultural level, to develop an ability to exchange signals that are new combinations of simple learned signals.

10. CONCLUSION

We have described a number of simulations in which neural networks are viewed in an Artificial Life perspective and we have applied this perspective to the study of language. As anticipated, the aspects of language that tend to be investigated using this perspective are different from the aspects of language that are typically studied within the symbol manipulation framework and also in classical connectionist simulations. The emphasis on the environment in which language evolves, develops, and is used; the consideration of all kinds of change that affect language in the individual and in populations of individuals (biological evolution, cultural evolution, biological development, learning); and the fact of viewing language as the concrete behavior of concrete organisms inevitably tend to shift the attention from the abstract study of phonological, morphological, and syntactical phenomena to the emergence of language from animal communication; the role of language in cognition; the relationship between language and sociality; and language as a populational rather than as an individual phenomenon.

The work described in the present paper undoubtedly represents only a first step toward an Artificial Life account of language. Furthermore, in many cases we have been forced to describe imaginary rather than actual simulations. Given the necessity to include ecologies, populations, genotypes, life histories, etc., in addition to the naked neural networks of classical connectionist simulations, it is clear that simulations with ALNNs will tend quickly to become very complex. It is also clear that the architectures of neural networks that must evolve and develop complex linguistic abilities for dealing with complex environments cannot remain forever the simple feedforward, three-layer architectures of most current connectionist simulations. In fact, one of the most important directions of research with ALNNs is the study of the coevolution of language and neural architecture and the evolutionary emergence of genetically transmitted predispositions to learn a language. We think that these problems can be solved and that even the preliminary work reported here shows that an Artificial Life approach to language is both feasible and exciting.

REFERENCES

- Cangelosi, A., Parisi, D., & Nolfi, S. 1994. Cell division and migration in a 'genotype' for neural networks. *Network*, **5**, 497–515.
- Denaro, D., & Parisi, D. 1995. *Cultural evolution in populations of neural networks*. Institute of Psychology, CNR, Rome.
- Diaz, R. M., & Berk, L. E. 1992. *Private speech. From social interaction to self-regulation*. Hillsdale, NJ: Erlbaum.
- Elman, J. L. 1990. Finding structure in time. *Cognitive Science*, **14**, 179–211.
- Elman, J. L. 1993. Learning and development in neural networks: The importance of starting small. *Cognition*, **48**, 71–99.
- Elman, J. L. 1995. Language processing. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks*. 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.
- Floreano, D., & Parisi, D. 1994. *A connectionist account of language development: imitation and naming*. Institute of Psychology, CNR, Rome.
- Fodor, J. A., & Pylyshyn, Z. 1988. Connectionism and cognitive architecture: A critical analysis. In S. Pinker & J. A. Fodor (Eds.), *Connections and symbols*. Cambridge, MA: MIT Press.
- Goldberg, D. E. 1989. *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.
- Hadley, R. F. 1994. Systematicity in connectionist language learning. *Mind and Language*, **9**, 247–272.
- Hare, M., & Elman, J. L. 1993. From weared to wore: A connectionist model of language change. In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Hamilton, W. D. 1964. The evolution of social behavior. *Journal of Theoretical Biology*, **7**, 1–52.
- Hinton, G. E., & Shallice, T. 1991. Lesioning an attractor network: investigation of acquired dyslexia. *Psychological Review*, **98**, 74–95.
- Holland, J. J. 1975. *Adaptation in natural and artificial systems*. Ann Arbor, Michigan: University of Michigan Press. (Reprinted by MIT Press, 1992.)
- 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.
- Jordan, M. I., & Rumelhart, D. E. 1992. Forward models: supervised learning with a distal teacher. *Cognitive Science*, **16**, 307–354.
- Langton, C. G. 1992. Artificial life. In L. Nadel & D. L. Stein (Eds.), *1991 Lectures in complex systems*. Reading, MA: Addison-Wesley.
- McClelland, J. L. 1989. Parallel distributed processing: implications for cognition and development. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology*. Oxford: Clarendon Press.
- McClelland, J. L., & Rumelhart, D. E. (Eds.) 1986. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and biological models*. Cambridge, MA: MIT Press.
- Miller, W. T., III, Sutton, R. S., & Werbos, P. J. (Eds.) 1990. *Neural networks for control*. Cambridge, MA: MIT Press.
- Mundinger, P. C. 1980. Animal cultures and a general theory of cultural evolution. *Ethology and Sociobiology*, **1**, 183–223.
- Nolfi, S., & Parisi, D. 1995a. Evolving artificial neural networks that develop in time. In F. Moran, A. Moreno, J. J. Merelo, & P. Chacon (Eds.), *Advances in artificial life*. New York: Springer.
- Nolfi, S., & Parisi, D. 1995b. “Genotypes” for neural networks. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks*. Cambridge, MA: MIT Press.
- Parisi, D., Cecconi, F., & Cerini, A. 1995. Kin-directed altruism and attachment behavior in an evolving population of neural networks. In N. Gilbert & R. Conte (Eds.), *Artificial societies: The computer simulation of social life*. London: UCL.
- Patterson, K., Seidenberg, M. S., & McClelland, J. L. 1989. Connections and disconnections: acquired dyslexia in a computational model of reading processes. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology*. Oxford: Clarendon Press.
- Plaut D. C., & Shallice, T. 1993. Deep dyslexia: a case study of connectionist neuropsychology. *Cognitive Neuropsychology*, **10**, 337–500.
- Plunkett, K., & Marchmann, V. 1991. U-shaped learning and frequency effects in a multi-layered perceptron: implications for child language acquisition. *Cognition*, **38**, 73–193.

- Plunkett, K., & Marchmann, V. 1993. From rote learning to system building. *Cognition*, **48**, 21–69.
- Reilly, R. G., & Sharkey, N. E. (Eds.) 1992. *Connectionist approaches to natural language processing*. Hillsdale, NJ: Erlbaum.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986. Learning internal representations by error propagation. In D. E. Rumelhart, & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations*. Cambridge, MA: MIT Press.
- Rumelhart, D. E., & McClelland, J. L. 1986. On learning the past tense of English verbs. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and biological models*. Cambridge, MA: MIT Press.
- Seidenberg, M. S., & McClelland, J. L. 1989. A distributed, developmental model of word recognition and naming. *Psychological Review*, **96**, 523–568.
- Sejnowski, T. J., & Rosenberg, C. R. 1987. Parallel networks that learn to pronounce English text. *Complex Systems*, **1**, 145–168.
- Stevens, K. N., & House, A. S. 1955. Development of a quantitative description of vowel articulation. *The Journal of the Acoustic Society of America*, **2**, 484–493.
- 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, MA: 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. *Modeling comprehension as behavior change*. Unpublished paper.