



ELSEVIER

Cognitive Development 17 (2002) 1301–1321

COGNITIVE
DEVELOPMENT

Artificial Life and Piaget

Domenico Parisi^{a,*}, Matthew Schlesinger^b

^a National Research Council, viale Marx 15, Rome 00137, Italy

^b Southern Illinois University, Carbondale, IL, USA

Abstract

Artificial Life is the study of all phenomena of the living world through their reproduction in artificial systems. We argue that Artificial Life models of evolution and development offer a new set of theoretical and methodological tools for investigating Piaget's ideas. The concept of an Artificial Life Neural Network (ALNN) is first introduced, and contrasted with the study of other recent approaches to modeling development. We then illustrate how several key elements of Piaget's theory of cognitive development (e.g., sensorimotor schemata, perception-action integration) can be investigated within the Artificial Life framework. We conclude by discussing possible new directions of Artificial Life research that will help to elaborate and extend Piaget's developmental framework. © 2002 Elsevier Science Inc. All rights reserved.

Keywords: Artificial Life; Piaget; Evolution

1. Introduction

Piaget cannot be said to be very popular these days among developmental psychologists and among psychologists in general. In any case, he is much less popular than he deserves to be given the great importance of his contribution to our understanding of behavior and cognition and their development in the child. (Of course there are exceptions; cf. Langer, 2000; Russell, 1996) The reason for Piaget's marginalization is not only that in the last decades developmental psychologists have discovered many significant facts about children's cognitive abilities that do not fit well within Piaget's theoretical framework—which is certainly true. A more fundamental reason seems to be that Piaget's theoretical framework makes

* Corresponding author. Tel.: +39-6860-90376; fax: +39-6824-737.

E-mail address: parisi@ip.rm.cnr.it (D. Parisi).

a number of fundamental assumptions which are opposite to the assumptions underlying much of current “cognitivist” psychology. We describe three of these assumptions.

The first assumption is Piaget’s “genetic epistemology,” that is, the idea that we can better (or only) understand X if we reconstruct how X has become what it is (Piaget, 1952). Hence, the central role assigned by Piaget to development for understanding intelligence and, more generally, to the study of how all aspects of behavior and cognition develop from birth on in order to understand adult behavior and cognition. This idea of a genetic epistemology has very little following among cognitivist psychologists. Cognitivist psychology is much closer to Chomsky’s idea that one must first understand an adult competence (e.g., the linguistic competence) and only then can one try to reconstruct using a model of the adult competence how this competence is acquired during ontogeny, than to Piaget’s position that an adult competence can only be understood if we first examine how that competence is progressively acquired during ontogeny. Some cognitivist psychologists even assume that not much develops at all since most of what is important in behavior already exists at birth, i.e., it is innate (e.g., Spelke, 1998; Wynn, 1992).

A second aspect of Piaget’s psychology which is quite alien to current cognitivist psychology is its fundamental biological orientation. Piaget conceives of intelligence as a form of adaptation which rests on fundamental principles analogous to those of biological adaptation, as an extension at a more advanced level of mechanisms and processes that underlie all biological phenomena (Piaget, 1971). And he is the author of a somewhat forgotten paper which proposes a neural model of the interiorization of actions (Piaget, 1947). Cognitivist psychology not only has a view of the mind according to which the mind is analogous to the software of a computer and therefore should be studied quite apart from studying the nervous system (which is analogous to the hardware of the computer) but, notwithstanding many cognitivist psychologists’ innatist and “evolutionary” assumptions, is completely extraneous to the idea that cognition is a form of biological adaptation and should be studied as such.

Finally, a third fundamental difference between Piaget and cognitivist psychology concerns their respective positions with respect to language. Language is central for cognitivist psychology and for classical (i.e., computational) cognitive science. In the last 40 years language has been studied by a branch of psychology called psycholinguistics which has played a crucial historical role in the “cognitive revolution” of the 50s and 60s and in the emergence of cognitive science. Much of psycholinguistics has been the union of psychology and Chomskian linguistics, which is especially concerned with keeping language separate from the rest of cognition. More fundamentally, cognitivist psychology being based on the analogy between mind and computer software tends to view all mental activity as symbol processing. For Piaget, on the contrary, language is only one aspect of cognition, is derived from non-linguistic (sensory–motor) cognition, and intelligence is fundamentally non-linguistic and is derived from action, not from language.

These three reasons (but there may be more) explain why Piaget is being largely ignored by contemporary cognitivist psychology. However, there are other, more intrinsic, reasons that play a role in explaining the present isolation and marginalization of Piaget and the sense of an interrupted growth which one feels when considering Piaget's work today. These additional reasons have to do with some critical limitations of his work. When one examines Piaget's theoretical constructions and empirical analyses it is difficult to avoid the impression that he has been able to obtain all the results that could have been obtained given the theoretical and methodological instruments that were available to him. But these instruments have intrinsic limitations. To further pursue Piaget's goals today and to continue his work while still adhering to his fundamental orientations it is necessary to use new theoretical and methodological tools that go beyond those that he used.

We believe that these new theoretical and methodological tools can be provided by Artificial Life. With respect to all three general issues we have discussed, Artificial Life takes Piaget's side rather than cognitivist psychology's side. First, Artificial Life shares Piaget's "genetic" view that in order to understand any phenomenon one has first of all to study how the phenomenon has become what it is. This is why Artificial Life is centrally concerned with evolution, development, growth, change. Second, as clearly indicated by its name, Artificial Life has a fundamentally biological orientation and, if concerned with behavior and cognition, cannot but view, like Piaget, behavior and cognition as basically biological phenomena. Third, for Artificial Life language is but one aspect of behavior and cognition, although quite important for characterizing the specific adaptive pattern of *Homo sapiens*. But Artificial Life's goal is not to explain the mind in terms of language but how language emerges both phylogenetically and ontogenetically in nonlinguistic organisms. More generally, Artificial Life uses neural networks to model behavior and cognition, which rules out any view of mind as basically constituted by symbol processing.

But what is even more important is that Artificial Life can provide new theoretical and methodological tools that are needed to go beyond Piaget's results while at the same time pursuing many of the same scientific goals as Piaget's. These new tools are (a) simulations as a new way of expressing scientific theories and models, (b) a set of specific modeling tools such as neural networks, genetic algorithms, and other Artificial Life models, and (c) the general theoretical framework of complex systems. We believe that using these new tools Piaget's work can re-acquire the centrality it deserves in the study of behavior and cognition and, what is more important, can be extended beyond Piaget's own results. Of course, pursuing Piaget's scientific goals and inserting his theoretical orientations within an Artificial Life framework will change in important ways many aspects of Piaget's work. But this of course is how scientific research should proceed.

The paper is articulated in the following way. In [Section 2](#), we briefly present Artificial Life. In [Section 3](#), we describe some aspects of Artificial Life simulations that have a Piagetian flavor. In [Section 4](#), we draw some conclusions.

2. Artificial Life

Artificial Life is the study of all phenomena of the living world through their reproduction in artificial systems (Langton, 1995). This ordinarily means to simulate living phenomena in a computer although sometimes physical artifacts (robots) may be constructed that exhibit some of the behaviors of real organisms (Nolfi & Floreano, 2000). Simulations are a new way of expressing scientific theories and hypotheses about the causes, mechanisms, and processes that underlie observed phenomena and explain those phenomena. Traditionally, scientific theories and hypotheses are expressed using the symbols of natural language and/or those of logic and mathematics. Simulations are scientific theories expressed as computer programs. When the program runs in the computer, the simulation results are the empirical predictions drawn from the theory. These results/predictions can then be compared with observed data to confirm or disconfirm the simulation/theory.

Simulations offer a number of important advantages as research tools. If one expresses one's theory or hypothesis in the form of a computer program, one is forced to be explicit, complete, detailed, because otherwise the program won't run in the computer or the simulation won't produce the expected results. Furthermore, theories expressed as simulations tend to have explicit, detailed, and rich empirical content since, as we have observed, the simulation results are the empirical predictions which are derived from the theory. Another advantage of simulations is they function as virtual experimental laboratories in which, as in the physical experimental laboratory, the researcher observes the phenomena in controlled conditions, manipulates the variables and parameter values that influence the phenomena, and determines the consequences of his or her manipulations. Finally, given the great memory and computing power of computers, one can simulate more complicated phenomena than those that can be observed in the physical laboratory, or phenomena that cannot be brought in the physical laboratory because they are long past or are too big or last too long. For example, one can reproduce the ecological conditions in which behavior is exhibited, without de-contextualizing the behavior as in real laboratory experiments, or one can simulate long-term evolutionary change in a population of organisms together with more short-term developmental change in the lifetime of the individual.

Artificial Life simulations address all sorts of phenomena of the living world, including the behavior, cognitive abilities, and mental life of organisms. A central characteristic of the behavior and cognitive abilities of organisms is that, like all other properties of organisms, they develop, that is, they change during an organism's life or, more specifically, during developmental age from birth (or conception) to maturity. Hence, Artificial Life can be used to study the same phenomena that were of interest to Piaget.

The behavior of organisms is controlled by their nervous system, although other systems of the body such as the endocrine system also play an important role. To

study the behavior of organisms Artificial Life uses (artificial) neural networks, which are simulation models directly inspired by the physical structure and way of functioning of the nervous system. A neural network is made up of units (neurons) that influence each other through unidirectional excitatory or inhibitory connections (synapses between neurons). A typical neural network has input (sensory) units connected to internal units connected to output (motor) units. Units have quantitative activation states (firing rate of neurons) and connections have quantitative weights (number of synaptic sites between neurons). The activation state of input units depends on physical and chemical events taking place outside the neural network whereas the activation state of internal and output units depends on the weighted excitations and inhibitions arriving to a unit from other connected units. The activation state of output units determines effects outside the neural network such as muscle contractions and therefore movements of the organism's body parts.

In psychology the name of the approach which uses neural networks to study the behavior and cognition of organisms is connectionism. However, when neural networks are used as part of the broader enterprise of Artificial Life (Parisi, 2001; Parisi, Ceconi, & Nolfi, 1990) they are somewhat different from the neural networks of classical connectionism (Rumelhart & McClelland, 1986). The main differences are that Artificial Life Neural Networks (ALNNs):

- (1) have a physical body
- (2) live in a physical environment
- (3) have an inherited genotype
- (4) are members of evolving populations.

2.1. ALNNs have a physical body

In classical connectionism neural networks tend to be viewed as abstract input/output devices or information processing systems. In Artificial Life neural networks are models of a physical organ (the nervous system) which is contained in a physical body which has a specific shape, size, physical disposition of sensory and motor systems, and other internal organs.

2.2. ALNNs live in a physical environment

Classical neural networks live in a void, or better, in an environment which coincides with the researcher who uses them. The researcher decides what is the network's input and evaluates the network's output. ALNNs live in and interact with a physical environment. It is the environment that provides the network with input and evaluates the network's output and, most important, the network can modify the external environment with its (motor) output and therefore it can influence the input which it receives from the environment (Nolfi & Parisi, 1993; Parisi, 1997a).

2.3. ALNNs have an inherited genotype

The body of an Artificial Life organism contains not only a neural network but also a genotype which each individual organism inherits from its parents. The genotype at least partially determines the individual's neural network and therefore the individual's behavior and cognitive activities.

2.4. ALNNs are members of evolving populations

Unlike classical connectionism Artificial Life does not study individual networks but always populations of individually different neural networks. An ALNN is born, develops, possibly reproduces, and dies. It is a member of a population of individuals which evolves because the individuals reproduce selectively and with the constant addition of new variants.

To study evolutionary change Artificial Life uses genetic algorithms. Populations of ALNNs live in an environment and reproduce selectively, based on the superior adaptation of some individuals and other factors (including chance). Reproduction means that new individuals are created that inherit the genotype of their parents. However, an offspring's genotype tends to be somewhat different from its parents' genotypes because of sexual recombination and random genetic mutations. Since an individual's genotype in part determines the individual's neural network, offspring behave similarly but not identically to the parents. Selective reproduction and the constant addition of new variants to the genetic pool results in evolutionary change across generations in the behavior of organisms.

As we have already observed, Artificial Life shares with Piaget three fundamental assumptions: a "genetic" view of behavior and cognition, a broadly biological orientation, and a view of nonlinguistic cognition as more basic than language. However, Artificial Life also differs from Piaget with respect to the first two of these three assumptions.

Both Artificial Life and Piaget share a "genetic" view of behavior: if one wants to understand the behavior of organisms, one should first reconstruct how the behavior has become what it is. For Piaget this means to study behavioral and cognitive development as a necessary step to understanding adult behavior and cognition. Artificial Life extends Piaget's genetic approach by including behavioral and cognitive evolution, that is, the study of how behavioral and cognitive abilities evolve in successive generations in a population of individuals. This has the important consequence that, since individuals inherit a genotype from their parents which represents the current (individual) outcome of past evolution in the population, one can study (simulate) how genetically inherited information and information extracted by the individual from its specific experience in the environment interact to produce the individual's behavioral phenotype and its development during the individual's life.

Both Artificial Life and Piaget share a broadly biological orientation. Behavior and cognition (or "intelligence") are both forms of biological adaptation. However,

Artificial Life goes further in its biological view of behavior than Piaget was probably ready to accept. This is not so much due to the fact, which we have already discussed, that Artificial Life studies evolutionary change together with developmental change, whereas Piaget was apparently convinced that psychology's task was to study the individual organism and how it develops, while evolutionary change in populations of organisms was to be entrusted to other disciplines such as evolutionary biology and genetics. What is more critical is that Artificial Life shares the connectionist claim that behavior and cognition are best studied using theoretical models explicitly inspired by the physical structure and way of functioning of the nervous system, i.e., neural networks. Piaget was quite explicit that the biological adaptation represented by intelligence was a "functional" level quite different from the "physical" level of the brain and of more elementary forms of biological adaptation (even if developed from these more elementary forms, e.g., Piaget, 1972). In both cases, Artificial Life appears to be less bounded by disciplinary separations while Piaget was more traditionally a psychologist.

3. Other simulation approaches to the study of development

In the last one or two decades development has been studied using computer simulations based on various theoretical models. Many but not all of these models use neural networks. We briefly describe these models and contrast them with the Artificial Life approach to development.

3.1. *Backpropagation of error*

The majority of connectionist models of development employ a feedforward neural network that is trained by backpropagation of error. Typically, "backprop nets" are presented with a fixed set of inputs, and the network's connection weights are gradually adjusted until the outputs of the network are sufficiently close to their desired values. Some of the phenomena that have been investigated using this approach include the development of perceptual categorization (Mareschal & French, 2000; Quinn & Johnson, 1997) and object-oriented behaviors in infants (e.g., reaching and visual tracking in Mareschal, Plunkett, & Harris, 1999; Munakata, McClelland, Johnson, & Siegler, 1997), language acquisition (Plunkett & Marchman, 1991), as well as the development of children's reasoning on the balance scale (McClelland, 1989).

A related set of models have also been proposed by Shultz and colleagues to simulate children's reasoning on several of Piaget's concrete and formal operational tasks, including the balance scale, conservation of number, and seriation (Mareschal & Shultz, 1993; Shultz, 1998; Shultz, Mareschal, & Schmidt, 1994). These models employ a combination of two learning algorithms—quickprop, a variant of backprop, and cascade correlation. An unusual and noteworthy feature of the cascade correlation algorithm is that it enables the network to alter its own

internal structure (i.e., by recruiting new hidden units), effectively increasing the representational power of the network.

3.2. *Adaptive resonance theory*

Despite the popularity of the backprop algorithm, several researchers have raised a number of concerns with its use for simulating cognitive development. For example, backprop may not adequately capture discontinuous or stagewise changes in children's behavior, a central tenet of Piaget's theory (e.g., Hartelman, van der Maas, & Molenaar; 1998; Raijmakers, van Koten, & Molenaar, 1996). As an alternative approach, Molenaar and colleagues have suggested the use of adaptive resonance theory (or ART, see [Carpenter, Grossberg, & Reynolds, 1991](#)) for modeling cognitive change.

ART models differ from conventional feedforward neural networks in at least two major ways. First, they are often represented by a system of differential equations, rather than a chain of feedforward processing units (e.g., see [Raijmakers & Molenaar, 1997](#)). Second, because ART networks are nonlinear dynamical systems, their internal patterns of activation persist and evolve over time. Like the models studied in dynamic field theory (see below), the internal activity of an ART network may exhibit a variety of dynamical properties, including periodic rhythms, resistance to external perturbation, and resonance across components of the system.

From a developmental perspective, there are several important features offered by ART: (1) greater biological plausibility than backprop, (2) simulation of continuity and discontinuity in development, and (3) like cascade correlation, a framework for modeling not only local (i.e., synaptic) but also global changes in network architecture.

3.3. *Hebbian nets*

A second alternative to backprop is Hebbian learning. Unlike backprop, Hebbian nets employ an unsupervised learning procedure. Learning occurs through gradual increases in the connections between units that fire synchronously (or alternatively, a decrease in connection strength during asynchronous firing). As an example of this approach, Munakata (1998) recently implemented a Hebbian net to simulate the development of the A-not-B error in young infants on Piaget's object permanence test. An important insight offered by Munakata's model is that perseverative responding in infants may be influenced by the interaction between long-term and short-term (i.e., latent and active) memory traces.

3.4. *Reinforcement learning*

A third alternative to backprop implements the principles of reinforcement learning (RL) as a framework for studying learning and development (for a

comprehensive introduction, see Sutton & Barto, 1998). Three key ideas from this approach are: (1) the organism is an autonomous agent that explores and interacts with its environment; (2) the organism behaves rationally with respect to achieving its goals, and (3) the environmental consequences of the organism's behavior (e.g., positive and negative reinforcement) drive the learning process.

When implemented as a learning rule in neural networks, RL simulates the process of variation-and-selection or trial-and-error learning. In contrast to backprop, RL agents are not explicitly taught a specific pattern of behavior. Instead, their spontaneous behaviors are evaluated with respect to a goal (much like the fitness formula used to evaluate ALNNs); adaptive behaviors are maintained and improved while other behaviors are discarded. Such models are ideally suited for simulating sensorimotor development, especially during infancy. For example, [Schlesinger and Parisi \(2001\)](#) use RL to simulate the development of visual tracking in infants. Consistent with Piaget's notion of sensorimotor cognition, this eye-movement model learns to generate a variety of future-oriented or prospective behaviors (e.g., anticipate the reappearance of an occluded object) using body-based or sensorimotor strategies.

In addition to these classical connectionist models, there are other modeling approaches that implement other types of architectures to investigate learning and development. We briefly describe three relevant approaches.

3.5. *Dynamic field theory*

Like ART, dynamic field theory (DFT) exploits the computational properties of nonlinear dynamical systems (e.g., [Thelen & Smith, 1994](#)). The dynamic field is a mathematical structure which simulates the interaction between global inhibition and local excitation in neural processing systems (though, properly speaking, DFT models are not neural networks). A recent example of a DFT model for studying development is offered by [Thelen, Schöner, Scheier, and Smith \(2001\)](#), who propose an alternative to Munakata's Hebbian network for modeling perseverative reaching in Piaget's A-not-B task.

3.6. *ACT-R*

In sharp contrast to neural network models, symbolic models such as ACT-R ([Andersen, 1993](#)) implement many of the key assumptions of the cognitivist approach (e.g., knowledge as symbolic representations, behavior as condition-action rules, etc.). Two recent examples of ACT-R models for studying development include simulations of numerical cognition in infants ([Simon, 1998](#)) and problem-solving in young children ([Jones, Ritter, & Wood, 2000](#)). While symbolic modeling is certainly incompatible with Piaget's notion of sensorimotor knowledge in infants, it should be noted that the ACT-R paradigm may complement Piaget's conceptualization of operational reasoning as a symbolic process.

3.7. Bayesian nets

A third alternative to connectionist models is an emerging approach that highlights statistical learning processes. Bayesian learning is a computational paradigm for identifying statistical regularities among inputs. This approach is particularly relevant for developmental researchers who investigate language acquisition and categorization development, where statistical dependencies play a central role. For example, Tenenbaum and Xu (2000) propose a Bayesian model of novel word acquisition in young children.

While all of the modeling approaches described thus far share the goal of investigating developmental processes, they differ from Artificial Life simulations in a number of important ways. First, only two of the models (Jones et al., 2000; Schlesinger & Parisi, 2001) simulate embodied organisms. Second, only these same two models simulate an environment with which the organism interacts. Third, because all of these models focus on ontogenetic change in individual organisms, they are unable to simulate changes at the population level, or more generally, phylogenetic change in a species. In contrast, ALNNs appear to be a more versatile framework, by incorporating not only embodiment and environmental interaction, but also the capacity for both developmental and evolutionary timescales of change. ALNN simulations can use all kinds of connectionist models of learning, but learning in ALNNs occurs on the basis of pre-existing, genetically inherited information, which is the result of another process of acquisition at the population level, i.e., evolution. Development, as distinct from learning, appears to be changes during life that depend both on this genetically inherited information and on experience and learning.

4. Examples of Artificial Life simulations with a Piagetian flavor

4.1. ALNNs as models of Piaget's sensory–motor schemata

ALNNs are basically sensory–motor neural networks. Their input encodes physico-chemical events and processes taking place in the external environment and their output encodes motor actions that have an influence on the external environment. However, unlike behaviorist stimulus–response systems and cognitivist condition–action rules, which directly map an input into an output, the input–output mapping in neural networks is mediated by a distributed processing system which confers to ALNNs some of the properties of Piaget's sensory–motor schemata. In particular, ALNNs are adaptive, that is, they possess both properties that Piaget considers as basic for intelligence: assimilation and accommodation. They are internally designed to both assimilate new inputs to their already existing internal structure and to accommodate (change) this internal structure as a function of new inputs (for a related interpretation, see Shultz, Schmidt, Buckingham, & Mareschal, 1995).

When a neural network is exposed to some input it has never experienced before, the neural network will always generate some output which in many cases will also be a plausible and reasonable output. This behavior is completely spontaneous and automatic: the neural network cannot but generalize to new inputs. It is the distributed nature of the ‘association’ that a neural network establishes between inputs and outputs which explains why a neural network tends to respond to new inputs, never experience before, with more or less plausible outputs by assimilating them through its already existing distributed network of connections and connection weights. Hence, unlike behaviorist stimulus–response devices and cognitivist condition–action rules, neural networks share the structural and systemic character of Piaget’s notion of a schema. Inside a neural network it is impossible to trace individual input/output associative links but it is the ensemble of all connections and connection weights which constitutes the collective support for all input/output associations, both already experienced and new. Therefore, a new input will never find a neural network unable to react appropriately. The new input will necessarily be assimilated to the neural network’s schemata, i.e., pattern of connection and connection weights. It will evoke an activation pattern in some set of internal units (an internal representation) and on the basis of this internal representation some plausible output will be generated.

In addition to being necessarily assimilatory, a neural network also accommodates. By accommodation Piaget means that already existing schemata always change as a function of the new experiences (inputs and input/output mappings) of the organism. As is well known, connectionist models are essentially learning models. Neural networks are not designed by the researcher. The researcher limits himself or herself to creating the initial conditions for learning and self-organization to occur. Therefore, when a neural network encounters a new experience it tends to change on the basis of the new experience. Learning in neural networks consists in modifications in the neural network’s connection weights or in the network’s architecture (pattern of pre-existing connections) and it can be supervised or unsupervised. Hence, through learning the neural network accommodates its internal structure of schemata. These accommodations change the internal representations that the neural network creates of all inputs, both old and new, and therefore how the network responds to all inputs.

4.2. *ALNNs integrate perception and action*

4.2.1. *Co-adaptation of sensory and motor capacities*

As we have said, ALNNs are basically sensory–motor systems. The input is a sensory input and the output is a motor output. However, again unlike behaviorist stimulus–response systems and cognitivist condition–action rules, in the internal structure of a neural network the sensory component and the motor component are not simply juxtaposed and associated but they are fused and integrated together. This integration of perception and action appears to be very much in the spirit of Piaget.

This Piagetian characteristic of ALNNs can be clarified by the following examples, in which the sensory and motor capacities of an organism constrain each other. Let us imagine a simple organism which lives in an environment with randomly distributed food elements. If the organism can only move ahead at fixed speed or turn to the left or to the right, it is necessary for the organism's neural network to receive inputs encoding the direction in which some currently perceived food element lies, but it is useless to know the distance of the food element. The problem that the organism must solve is to turn in such a way that its direction of movement coincides with the direction in which the food element lies, and then move ahead. To be informed by the senses also about the distance of the food element is useless or, worse, disrupts the organism's neural network which would have to process additional and useless information. In fact, with organisms with such a simple motor system to add sensory information about food distance leads to less efficient behavior. But consider another organism endowed with a two-segment arm which must reach objects with the tip of the arm (hand). For this organism it is essential to be informed by its senses about both the direction and the distance of the objects to be reached. To know only the direction in which an object lies would make it impossible for this organism to reach the object with its arm (e.g., Parisi et al., 1990).

These examples show that it is necessary for the sensory and motor capacities of an organism to be coordinated and co-adapted. For an organism to receive less or more sensory information than the organism's motor system can utilize leads to inefficient organisms. Hence, we would expect that in simulations in which sensory and motor capacities evolve in the population or develop in the individual, they co-evolve or constrain each other in the course of development in such a way that they are constantly more or less co-adapted.

4.2.2. The importance of actions in structuring the internal organization of neural networks

Basically, neural networks are devices that transform activation patterns into other activation patterns. A sensory–motor neural network is a network that receives a sensory activation pattern in its input units and transforms this sensory activation pattern into an appropriate motor pattern in its output units. This transformation takes place through a succession of intermediate activation patterns that can be observed in the successive layers of the internal units of the neural network's architecture. The weights of the connections linking each layer to the successive layer are the causal agents that effect these transformations.

Given that this is the basic task of a sensory–motor neural network, it is necessary that the network's connection weights which realize the successive transformations are such that they obey the following two general principles:

- (1) Make the activation patterns originating in inputs that must be responded to with the same action more similar to each other,
- (2) Make the activation patterns originating in inputs that must be responded to with different actions more different from each other.

Activation patterns can be represented as single points in the abstract hyperspace that corresponds to the layer of units in which the activation pattern is observed. Hence, each layer of units has its own hyperspace. The hyperspace corresponding to some particular layer of units has as many dimensions as there are units in the layer and the point which represents a particular activation pattern observed in the layer is located in a position that, for each dimension of the hyperspace, corresponds to the activation level of the unit representing that dimension. Given this type of representation, similar activation patterns in a layer of units will be represented by points that lie close to each other in the hyperspace, whereas different activation patterns will be represented by points that are more distant in the hyperspace.

If we give the name of “cloud” to the set of points that represent all the activation patterns evoked by inputs that must be responded to with the same action on the part of the organism, we see that learning or development consists in progressively modifying the network’s connection weights in such a way that, in the successive layers of units of the network’s architecture, the size of individual clouds becomes smaller and the distance between the centers of pairs of different clouds become greater. With progressively smaller clouds and, therefore, more similar internal representations of inputs which must be responded to with the same action, it becomes easier for the network to associate the same output activation pattern (action) to the different inputs represented in the cloud. And vice versa, with increasing distance among different clouds, which means more different internal representations of inputs that must be responded to with different actions, it becomes easier for the network to discriminate among inputs which must be responded to with different output patterns (actions).

This clearly shows that internal representations in sensory–motor networks are dictated by the motor actions with which the organism must respond to the various sensory inputs, rather than by those sensory inputs themselves. This appears to be very consonant with Piaget’s constructivist emphasis on the importance of action in the development of cognition.

4.3. Circular reactions: predicting the results of one’s own actions

A crucial property of ALNNs is that they are ecological networks, i.e., networks that live in a physical environment. What is simulated in an Artificial Life simulation is not only a neural network and what takes place inside the network but also the neural network’s physical environment and what takes place in the physical environment as a consequence of the neural network’s behavior. Living in a physical environment implies that the network’s input is determined by the environment and the network’s output causes changes either in the physical relation of the organism’s body or body part (eye, face, arm, etc.) to the environment or in the physical environment itself (objects are displaced by the organism, modified, taken apart, put together, etc.). In both cases the network’s input from the environment is influenced by the organism’s previous action and, therefore, unlike

classical connectionist networks, ALNNs at least partly determine their own input (Parisi, 1997a).

This influence of ALNNs on their environment and, as a consequence, on their own input is crucial for understanding the behavior of organisms. In many cases organisms behave in ways that are only aimed at exposing themselves to useful inputs, i.e., inputs that the organism knows how to respond to, or from which the organism can learn more than from other inputs, or simply inputs that can be used by the organism to confirm its predictions concerning the consequences of its actions (see below). Therefore, to understand the behavior of organisms it is equally important to reconstruct the network-internal cause–effect chain that leads from sensory input to motor output and the network-external cause–effect chain that leads from motor output to sensory input. These are Piaget’s circular reactions.

An important consequence of circular reactions in ALNNs is that ALNNs can learn to predict the consequences of their own actions. Each input–output cycle includes two steps. In step 1, the network’s sensory input produces an output activation pattern which encodes a motor action. However, the motor action is not physically executed but the output activation pattern is used as a self-generated input which re-enters the network. In step 2, the network responds to this self-generated input together with the current input from the environment to generate another output activation pattern which encodes a prediction of the next input from the environment that will result when the action will be actually executed. At this point the motor action is physically executed and its actual consequences are compared with the predicted consequences, using the backpropagation procedure to change the network’s connection weights so that the network learns to generate progressively more accurate predictions (Parisi et al., 1990).

Much behavior in the very first months of life in human infants can be interpreted as behavior which is generated with the sole purpose of learning to predict its consequences using the two-step procedure we have described. Initially, the infant learns to predict the changes in the inputs that originate from its own body (primary circular reactions). Then, the infant learns to predict changes in inputs originating from objects different from its own body, due to changes caused by the infant’s behavior in the physical relation of its body or body parts to these objects or in the objects themselves (secondary reactions). Finally, through systematic variations in behavior, the infant learns to organize and coordinate these predictions (tertiary circular reactions).

4.4. *Evolution and development*

Given his general biological orientation Piaget was clearly convinced of the importance of evolution and genetic inheritance but he apparently attributed a direct role to genetically inherited information only with respect to the very beginning of development (e.g., reflexes) and preferred to interpret most development as the result of self-organization and internal re-structuring (Piaget, 1952, 1954). In any case, given the methodological tools that were available to him, he was never

able to propose and test explicit hypotheses on what is genetically inherited and how genetically inherited information interacts with the organism's experience to determine developmental changes.

Artificial Life is as much concerned with evolution in populations of organisms as with development and learning in the individual organism. In Artificial Life simulations genetic information is explicitly encoded in genotypes (for example, as bitstrings) and this information is transmitted from parents to offspring. Since reproduction is selective, i.e., some individuals generate more offspring than others, and offspring's genotypes are similar but not identical to parents' genotypes because of sexual recombination (of portions of the mother's genotype with portions of the father's genotype) and because of genetic mutations that change randomly chosen parts of the inherited information, a population's genotypes change across successive generations in ways that make the population more adapted to the particular environment in which the population happens to live. (But a number of additional factors such as noise in the selection process, internal constraints such as pre-adaptations and exaptations, and simple chance can make the evolutionary process less than perfectly adaptive; see [Miglino, Nolfi, & Parisi, 1996](#)).

All kinds of information may be encoded in genotypes, e.g., information about the body of the organism (size, shape, etc.), the physical disposition of the organism's sensory receptors and movable body parts, and the organism's neural network (nervous system). The genotype can specify what is encoded in the network's input units and how, what is the neural network's pattern of connectivity (the network architecture), the network's connection weights, and other properties such as the parameters governing how the network learns as a function of experience. All information encoded in the genotype varies from one individual to another individual and is subject to change in the population during evolution.

The particular genetic information inherited by an individual plays a role in determining the individual's neural network and therefore the individual's behavior. (Notice, however, that an individual's behavior depends both on the individual's neural network and on the particular environmental encounters of the individual.) However, genetic information can play a more or less important role in determining the individual's neural network and it can find its phenotypical expression at different times during the lifetime of the individual. For example, in some Artificial Life simulations the network architecture is identical in all individuals and is fixed (i.e., decided once and for all by the researcher) and the genotype directly encodes the network's connection weights. The connection weights do not change at all during the individual's life, which means that there is no learning and no development. The individual's behavior is entirely innate and the individual's responds in the same to any given input throughout its entire lifetime. These simulations may more or less realistically simulate the behavior of very simple organisms (e.g., insects) but are clearly inadequate to model more complex organisms such as humans. However, this type of simulation can be used to study developmental processes (e.g., developmental stages) at a more abstract level by interpreting the genetic algorithm not as a model of evolution but as an abstract model of individual

development. For example, evolutionary changes at the population level, i.e., in successive generations of individuals, can tell us something about developmental changes in the individual (Schlesinger, Parisi, & Langer, 2000). (A more realistic use of the genetic algorithm as a model of individual development would be possible if one interprets neural development in the individual as a competitive process of selection among different groups of neurons; cf. Edelman's Neural Darwinism, in Edelman, 1987.)

However, if one interprets simulated genotypes as encoding genetically inherited information and the genetic algorithm as a model of evolution, one has to confront two problems: how and when genetic information is mapped into phenotypical information and how genetic information interacts with information from the individual's experience in the environment to determine the phenotype and its development (Parisi, 1996).

In many Artificial Life simulations information encoded in the genotype is directly mapped (one-to-one) into phenotypical information. For example, the quantitative value of each individual connection weight is directly encoded in a separate portion of the genotype as a real or binary number. Furthermore, the mapping is all completely executed at birth in such a way that, unless there is some role of experience and learning in changing the individual, the individual does not change and does not develop at all during life. However, in other simulations the phenotype/genotype mapping is more complex and it may not be instantaneous but what is actually encoded in the genotype is a developmental (or maturational) program. The genotype specifies not only what properties the phenotype must possess but also at what age of the individual these properties become manifest (Cangelosi, Parisi, & Nolfi, 1994).

Behavioral changes during an individual's life in minimally complex organisms are a function of both the information contained in the inherited genotype and the information derived from the individual's experience in the environment. In some Artificial Life simulations the connection weights of an individual's neural network change during an individual's life as the individual learns from its experience, i.e., the input and the teaching input received from the environment (backpropagation learning). However, the initial values of the connection weights or the learning parameters (learning rate and momentum), which both play a significant role in backpropagation learning, are genetically inherited and evolve in such a way that the individuals of later generations are able to learn more from their experience than the individuals of early generations because evolution finds more appropriate values for the initial connection weights and for the learning parameters (Bewley, McInerney, & Schraudolph, 1991). In other simulations the network architecture is encoded in the genotype and it evolves but the connection weights are not inherited; they are randomly assigned at birth to each individual and modified during the individual's life as the individual learns (Di Ferdinando, Calabretta, & Parisi, 2001). In still other simulations the network architecture evolves but "neural development" in the individual is influenced by the environment (Nolfi, Miglino, & Parisi, 1994; Nolfi & Parisi, 1996). In all these simulations, evolution

and learning/experience both determine the developmental changes that occur in the individual and they cooperate to equip the individual with the most appropriate behaviors (Belew & Mitchell, 1996).

To entrust evolution with the task of finding the most appropriate network architecture and learning with the task to find the most appropriate connection weights for this architecture appears to be a general strategy which is frequently adopted by nature (Elman et al., 1996). The individual learns during life starting with random initial connection weights (not inherited) but the individual's network architecture (or at least its general architectural schema) is genetically inherited and it is an architecture which has evolved in past generations of the population of which the individual is a member.

For the learning of many types of behaviors the most appropriate network architecture may be a modular one, with some connections (neural module) dedicated to one sub-task and other connections (other neural module) dedicated to another sub-task. For example, Rueckl, Cave, and Kosslyn (1989) have shown that networks with a modular architecture learn more easily than nonmodular ones the complex What and Where task that consists in two sub-tasks: given an object encoded in a "retina" as input, recognizing both what an object is (What sub-task) and where it is located (Where sub-task). Di Ferdinando et al. (2001) have shown that if the network architecture evolves in a succession of generations and is genetically inherited, the individuals of later generations tend to possess modular architectures rather than nonmodular architectures if the reproduction criterion (fitness) is how much an individual learns during life the complex task of recognizing both what an object is and where it is located. Since modular architectures learn the What and Where task more easily, it is these architectures that evolve in the population rather than nonmodular ones (Di Ferdinando et al., 2001).

On the other hand, if both the network architecture and the network's connection weights are encoded in the genotype, i.e., everything evolves and is genetically inherited and there is no learning during life, the results are less good. In these circumstances good genotypes, i.e., genotypes that result in the appropriate behavior, are hard for evolution to find because of a problem of "genetic linkage". Modular architectures appear to be better than nonmodular ones for learning the What and Where task because they are able to solve the problem of neural interference. In nonmodular architectures single connections are involved in both sub-tasks and therefore these connections may receive conflicting messages for changing their weight values (neural interference) during learning, which represents an obstacle to learning both sub-tasks well. On the other hand, in modular architectures this problem is solved because distinct connections are involved in the two sub-tasks and there is no possibility that connections receive conflicting messages. In genotypes encoding both modular architectures and their connection weights there may be a similar problem of genetic interference (or genetic linkage). In modular architectures the connection weights of distinct neural modules are encoded in distinct portions of the genotype. If a favorable genetic mutation happens in one such portion (encoding one module) while a nonfavorable one happens in another

portion (encoding another module), evolution can either keep (reproduce) the mutated genotype with a resulting advantage for the first sub-task but a disadvantage for the second sub-task, or drop the genotype, avoiding the disadvantage for the second sub-task but missing the advantage for the first one. There seems to be no way out of this dilemma. This may be why behaviorally complex organisms, that is, organisms that have to solve different sub-tasks while dealing with the same input, tend to inherit their modular network architecture but learn during their life the connection weights for the inherited architecture which are appropriate for exhibiting the appropriate behaviors (Calabretta, Di Ferdinando, Wagner, & Parisi, *in press*).

5. Conclusion

Artificial Life simulations are a useful tool for studying development in ways that may have pleased Piaget and for furthering some of Piaget's scientific goals. Artificial Life simulations use neural networks to model the behavior of organisms but, unlike classical connectionism, they simulate not only a neural network and what takes place inside the network but also the organism's body, its environment, its genotype, the interactions between the organism and the environment, and the evolutionary process at the population level which results in the organism's genotype. Basically, the neural networks used in Artificial Life simulations are sensory-motor systems. They receive sensory input from outside and map this input into motor actions. Given the distributed character of neural networks, inputs are mapped into outputs not in a one-to-one fashion but through an internal organization of the neural network (the network's architecture and connection weights) which resembles Piaget's notion of a schema. A schema is an adaptive device which automatically assimilates new inputs to already existing schemata and modifies (accommodates) existing schemata as a function of new inputs. The network's internal organization with the representations it generates (i.e., the activation patterns observed in sub-sets of the network's internal units) tends to make dissimilar inputs more similar if the inputs must be responded to with the same motor action, and similar input more dissimilar if they must be responded to with different motor actions. This implies that it is action rather than sensory input which dictates the form of internal representations and argues in favor of an action-based view of knowledge which appears to be quite Piagetian.

Since in Artificial Life simulations neural networks live in and interact with a physical environment, their motor output changes the environment and these changes are reflected in the network's sensory input. Hence, the network's reactions become circular reactions. The importance of circular reactions for the development of behavior, which was very much stressed by Piaget, is increased if the neural networks learn to be able to predict the consequences of their motor actions on the environment and therefore their future inputs.

Piaget agreed that development results both from inherited adaptations and from the experience and activity of the individual organism in its environment. However, he had no tools at his disposal for generating and testing explicit hypotheses for inherited adaptations. Artificial Life uses genetic algorithms as such tools. Simulated genotypes evolve in populations of organisms and they contain information which interacts with information from the individual's experience to produce the individual's behavior and development.

In conclusion, Artificial Life provides some important tools for studying development in a Piagetian perspective but most of the work is still to be done. Artificial Life neural networks so far have been mostly used to study basic sensory–motor adaptations but Piaget, while considering these adaptations as the basis on which later, more “cognitive,” adaptations are built, has contributed importantly to our understanding of these later, more “cognitive,” adaptations. For example, for Piaget at some stage of development physical actions become internal “operations.” The question is: can Artificial Life simulations deal with “operations?” We think that the answer is Yes (Parisi, 1997a, 1997b). Complex organisms such as humans have a mental life in addition to behavior. Sensory–motor neural networks basically have connections that go “forward” from sensory input to motor output. Mental life can be simulated using neural networks with a rich structure of connections that go “backward” and produce self-generated input for the neural network. The predictive networks described in Section 3.3 are an example of such networks. Motor actions are encoded in some sub-set of internal units but they are not physically executed. They re-enter the neural network as self-generated input which is used to produce predictions on their consequences. Self-generated input due to “backward” connections are the basis for all of mental life: mental images, remembering, thinking, reasoning, planning, etc. However, most of the work in this area is still a task for the future.

For Piaget much development is the result of changes due to internal re-organization triggered by experience. However, inherited information encoded in inherited genotypes clearly plays a role in development, and not only in the very earliest stages as Piaget sometimes seemed to think. Neural networks as complex systems with a distributed encoding of information can capture change as internal re-organization, but Artificial Life can also formulate and test explicit and specific hypotheses on genetically inherited information and simulate the interactions between this information and the information provided by the experience and activity of the individual organism that result in its development. However, what has been done so far mostly shows the feasibility of the approach and suggests some useful directions of research. In this area too, most work is still to be done.

References

- Andersen, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- Belew, R. K., & Mitchell, M. (Eds.). (1996). *Adaptive individuals in evolving populations*. Reading, MA: Addison-Wesley.

- Belew, R. K., McInerney, J., & Schraudolph, N. N. (1991). Evolving networks: Using the genetic algorithm with connectionist learning. In C. G. Langton et al. (Eds.), *Proceedings of the second conference on artificial life*. Reading, MA: Addison-Wesley.
- Calabretta, R., Di Ferdinando, A., Wagner, G., & Parisi, D. (in press). What does it take to evolve behaviorally complex organisms? *Biosystems*.
- Cangelosi, A., Parisi, D., & Nolfi, S. (1994). Cell division and migration in a 'genotype' for neural networks. *Network*, 5, 497–515.
- Carpenter, G. A., Grossberg, S., & Reynolds, J. H. (1991). ARTMAP: Supervised real-time learning and classification of nonstationary data by a self-organizing neural network. *Neural Networks*, 4, 565–588.
- Di Ferdinando, A., Calabretta, R., & Parisi, D. (2001). Evolving modular architectures for neural networks. In R. M. French & J. P. Sougné (Eds.), *Connectionist models of learning, development, and evolution*. London: Springer.
- Edelman, G. E. (1987). *Neural Darwinism. The theory of neuronal group selection*. New York: Basic Books.
- Elman, J. L., Bates, E., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness. A connectionist perspective on development*. Cambridge, MA: MIT Press.
- Jones, G., Ritter, F. E., & Wood, D. J. (2000). Using a cognitive architecture to examine what develops. *Psychological Science*, 11, 93–100.
- Langer, J. (2000). The descent of cognitive development. *Developmental Science*, 3, 361–378.
- Langton, C. G. (Ed.) (1995). *Artificial Life: An overview*. Cambridge, MA: MIT Press.
- Mareschal, D., & French, R. (2000). Mechanisms of categorization in infancy. *Infancy*, 1, 59–76.
- Mareschal, D., & Shultz, T. R. (1993). A connectionist model of the development of seriation. *Proceedings of the fifteenth annual conference of the cognitive science society* (pp. 170–175). Hillsdale, NJ: Lawrence Erlbaum.
- Mareschal, D., Plunkett, K., & Harris, P. (1999). A computational and neuropsychological account of object-oriented behaviours in infancy. *Developmental Science*, 2, 306–317.
- McClelland, J.L. (1989). Parallel distributed processing: Implications for cognition and development. In R. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 9–45). New York: Oxford University Press.
- Miglino, O., Nolfi, S., & Parisi, D. (1996). Discontinuity in evolution: How different levels of organization imply pre-adaptation. In R. K. Belew & M. Mitchell (Eds.), *Adaptive individuals in evolving populations* (pp. 399–415). Reading, MA: Addison-Wesley.
- Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics. The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
- Nolfi, S., & Parisi, D. (1993). Self-selection of input stimuli for improving performance. In G. A. Bekey & K. Y. Goldberg (Eds.), *Neural networks in robotics* (pp. 403–418). Boston: Kluwer.
- Nolfi, S., & Parisi, D. (1996). Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, 5, 75–98.
- Nolfi, S., Miglino, O., & Parisi, D. (1994). Phenotypic plasticity in evolving neural networks. In D. P. Gaussier & J. -D. Nicoud (Eds.), *From perception to action*. Los Alamitos, CA: IEEE Computer Society Press.
- Parisi, D. (1996). Computational models of developmental mechanisms. In R. Gelman & T. K. Au (Eds.), *Perceptual and cognitive development*. San Diego, CA: Academic Press.
- Parisi, D. (1997a). Active sampling in evolving neural networks. *Human Development*, 40, 320–324.
- Parisi, D. (1997). Artificial Life and higher level cognition. *Brain and Cognition*, 34, 160–184.
- Parisi, D. (2001) Neural networks and Artificial Life. In D. Amit & G. Parisi (Eds.) *Frontiers of life*. San Diego: Academic Press.
- Parisi, D., Cecconi, F., & Nolfi, S. (1990). Econets: neural networks that learn in an environment. *Network*, 1, 149–168.
- Piaget, J. (1947). Le problème neurologique de l'intériorisation des actions en opérations réversibles. *Archives de Psychologie*, 32, 241–258.
- Piaget, J. (1952). *The origins of intelligence in children*. New York: International Universities Press.

- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.
- Piaget, J. (1971). *Biology and knowledge. An essay on the relations between organic regulations and cognitive processes*. Chicago: University of Chicago Press.
- Piaget, J. (1972). *The psychology of intelligence*. Totowa, NJ: Littlefields Adams.
- Plunkett, K., & Marchman, V. (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition*, 38, 1–60.
- Quinn, P. C., & Johnson, M. H. (1997). The emergence of perceptual category representations in young infants: A connectionist analysis. *Journal of Experimental Child Psychology*, 66, 236–263.
- Raijmakers, M. E. J., & Molenaar, P. C. M. (1997). Exact ART: A complete implementation of an ART network. *Neural Networks*, 10, 649–669.
- Rueckl, J. G., Cave, K. R., & Kosslyn, S. M. (1989). Why are “what” and “where” processed by separate cortical visual systems? A computational investigation. *Journal of Cognitive Neuroscience*, 1, 171–186.
- Rumelhart, D. E. & McClelland, J. L. (1986). *Parallel distributed processing. explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.
- Russell, J. (1996). *Agency. Its Role in Mental Development*. Hillsdale, NJ: Erlbaum.
- Schlesinger, M., & Parisi, D. (2001). The agent-based approach: A new direction for computational models of development. *Developmental Review*, 21, 121–146.
- Schlesinger, M., Parisi, D., & Langer, J. (2000). Learning to reach by constraining the movement search space. *Developmental Science*, 3, 67–80.
- Shultz, T. R. (1998). A computational analysis of conservation. *Developmental Science*, 1, 103–126.
- Shultz, T. R., Mareschal, D., & Schmidt, W. C. (1994). Modeling cognitive development on balance scale phenomena. *Machine Learning*, 16, 57–86.
- Shultz, T. R., Schmidt, W. C., Buckingham, D., & Mareschal, D. (1995). Modeling cognitive development with a connectionist algorithm. In T. J. Simon & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling* (pp. 205–261). Hillsdale, NJ: Lawrence Erlbaum.
- Simon, T. J. (1998). Computational evidence for the foundations of numerical competence. *Developmental Science*, 1, 71–78.
- Spelke, E. S. (1998). Nativism, empiricism, and the origins of knowledge. *Infant Behavior and Development*, 21, 181–200.
- Tenenbaum, J. B., & Xu, F. (2000). Word learning as Bayesian inference. In L. R. Gleitman & A. K. Joshi (Eds.), *Proceedings of the 22nd annual conference of the cognitive science society* (pp. 512–517). Hillsdale, NJ: Lawrence Erlbaum.
- Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, 24, 1–34.
- Wynn, K. (1992). Addition and subtraction by human infants. *Nature*, 358, 749–750.