

In S.J. Scher and F. Rauscher (eds)
Evolutionary Psychology: Alternative Approaches.
Dordrecht: Kluwer, 2003, pp. 243-265

Evolutionary psychology and Artificial Life

Domenico Parisi

Institute of Psychology
National Research Council, Rome
parisi@ip.rm.cnr.it

1 Evolutionary psychology and neural networks

In this chapter we outline a research perspective which shares the basic goal of narrow evolutionary psychology, that is, to develop a “psychology informed by the fact that the inherited architecture of the human mind is the product of the evolutionary process” (Cosmides, Tooby, & Barkow, 1992, p. 7), but approaches this goal using a theoretical and methodological framework very different from that of narrow evolutionary psychology. There are two main characteristics which distinguish our approach from narrow evolutionary psychology. First, evolutionary psychologists tend to be cognitivists whereas the approach proposed here is connectionist. Cognitivism assumes that behavioral and mental processes can and should be studied at a functional or information-processing level without considering the physical structure of the brain. Cognitivists equate the mind with a computer’s software, which is analyzed and designed by computer scientists who ignore the physics of the machine. In contrast, connectionism is the idea that behavior and mental processes are best studied using theoretical models such as neural networks which are directly inspired by the physical structure of the nervous system. Second, practitioners of narrow evolutionary psychology use a variety of methods and data for developing and testing hypotheses on the inherited features of the human mind and on the selective pressures that have caused the evolutionary emergence of these features: experiments on human subjects, anthropological data concerning different cultures, various types of evidence on past human behavior and societies (Buss, 1999). The approach proposed here uses computer simulation as the main methodology for developing hypotheses on the evolution of the human mind and for deciding if a proposed theory actually predicts/explains the empirical evidence. Neural networks are simulations, i.e., theoretical

models expressed not using verbal or mathematical symbols, as traditionally in science, but as computer programs.

Why are neural networks better theoretical and methodological tools for studying the evolutionary shaping of the modern human mind than the cognitivist framework embraced by narrow evolutionary psychologists? One reason is theoretical. Neural networks complete the revolution in the study of the mind which began in the second half of the nineteenth century with the birth of scientific psychology. During the entire nineteenth century psychology has remained a halfway revolution because, while it has adopted the experimental and quantitative methods of the natural sciences, it still uses a mentalistic explanatory vocabulary of belief, goal, intention, representation, rule, etc., derived from common sense and from the philosophical tradition, which is different from the conceptual vocabulary of the natural sciences that recognizes only physical causes producing physical effects. Neural networks complete the revolution because what takes place in the interactions between a neural network and the external environment and inside the neural network is nothing but physical causes producing physical effects. Since evolutionary psychology aims at “conceptual integration” (Cosmides, Tooby, & Barkow, 1992, p. 4) between psychology and the biological sciences, it is easier to attain conceptual integration if the fields to be integrated speak the same language. Neural networks speak the same language of the neurosciences and, if they are viewed as a chapter of Artificial Life, as we will explain in the next section, they also speak the same language of the other biological sciences, in particular of evolutionary biology.

Notice however that, notwithstanding their conceptual unification with the natural sciences, neural networks are not reductionist if by reductionism we mean that facts about behavior and mental life become facts about neurons and synapses. Nervous systems are complex systems in which a very large number of neurons interact locally, and from these interactions emerge the global properties that we call behavior and mental life. Behavioral and mental phenomena are caused by neurons and synapses but cannot be predicted or deduced from even a perfect knowledge of neurons and synapses. Furthermore, in the case of humans behavior and mental life are caused as much by the cultural environment that humans themselves create as they are by neurons and synapses.

The other reason why neural networks are to be preferred for studying the evolved basis of the human mind is methodological. Given the dearth of direct empirical evidence for testing hypotheses on the adaptive bases of human traits, evolutionary psychology tends to be limited to analysis and argument and its adaptive explanations often seem to be “just so stories”. The causal history behind

many evolved human traits is very complicated (and not necessarily adaptive) and it is often impossible to demonstrate to the satisfaction of everyone that a verbally expressed theory actually explains the human trait it is intended to explain. The author of the theory may claim that his/her theory explains the evolutionary origin of language or cooperative behavior but other people may remain unconvinced, and there may be no way to decide the issue if one can only discuss verbally if the theory actually explains the facts, i.e., if the facts “derive” from the theory. Verbally expressed theories, such as those of narrow evolutionary psychology, are likely to be incomplete, insufficiently precise and detailed, and possibly internally inconsistent, so that no one can really say that some particular theory actually predicts or explains what it purports to explain. On the other hand, if a theory is expressed as a computer program - i.e., it is a simulation - when the simulation runs in the computer the results of the simulation are the predictions derived from the theory and these predictions can be seen by everyone (literally, on the computer’s screen). A theory expressed as a simulation must necessarily be complete, precise, detailed, and internally consistent because otherwise the program would not run in the computer or it would not generate the expected results. If the simulation results match the known empirical facts, the theory incorporated in the simulation actually explains the empirical facts. Furthermore, simulations function as virtual experimental laboratories in which the researcher can observe (simulated) phenomena in controlled conditions and can manipulate these conditions and determine the effects of the manipulations. This will strengthen our faith in the theory incorporated in the simulation. Hence, if one has a theory that claims to identify the evolutionary basis of some present human trait, one should incorporate his or her theory in a simulation and show that the simulation produces that human trait.

2 Neural networks as Artificial Life

Neural networks are the main research instrument of the approach to studying the evolutionary bases of human mind proposed in this chapter. However, the neural networks that are appropriate for investigating the evolutionary bases of the human mind are not the “classical” neural networks of Rumelhart and McClelland (1986). If one wants to study how “the inherited architecture of the human mind is the product of the evolutionary process” one has to view and use neural networks in the broader perspective of Artificial Life. Artificial Life is the attempt to understand all the phenomena of the living world by reproducing these phenomena in a computer. Nervous systems, behaviors, and mental phenomena are part of the living world and therefore neural networks are part of Artificial Life. However, neural networks explicitly considered as Artificial Life (Artificial Life

Neural Networks; ALNNs) are different from “classical” neural networks in the following respects (Parisi, Cecconi, & Nolfi, 1990):

- ALNNs have a physical body
- ALNNs live in a physical environment
- ALNNs have an inherited genotype
- ALNNs are members of biologically and, in the case of humans, also culturally evolving populations of networks.

A typical Artificial Life simulation aimed at behavioral and mental phenomena is a simulation of an entire population of organisms made up of successive generations of individuals. The population lives in an environment and the behavior of each individual in the environment is controlled by a neural network that simulates the individual’s nervous system. In addition to the neural network, each individual has a body with given physical properties (size, shape, spatial arrangement and physical properties of sensory and motor systems, internal organs and systems in addition to the nervous system) and an inherited genotype. Each individual is born, develops, possibly reproduces, and dies. The birth of an individual is the creation of a new organism which inherits the genotype of its two parents (if reproduction is sexual). The inherited genotype interacts with the external environment (external to the genotype and external to the organism’s body) to determine the phenotypical traits of the adult individual through a succession of developmental steps.

Evolution and genetic inheritance are intrinsic properties of Artificial Life simulations. Artificial Life uses genetic algorithms to model evolution. Individuals reproduce differentially, i.e., some individuals reproduce more than others, and reproduction is accompanied by the creation of new variants of genotypes due to random genetic mutations and to sexual recombination. Selective reproduction and the constant addition of new variants result in biological evolution, that is, changes in genotypes and, as a consequence, in phenotypes across a succession of generations. Inherited genotypes reflect the past history of adaptation to the environment of the species of which the individual organisms are members. Hence, ALNNs are ideal for studying the “inherited architecture” of the human mind which is the goal of evolutionary psychology. But actually doing simulations with even simple neural networks in simple environments shows that evolutionary change, and the inherited information which is its product, is not necessarily adaptive, as narrow evolutionary psychologists tend to think, but can result from non-adaptive causes such as chance, historical contingency, continuing a previous evolutionary path, and developmental constraints.

This is another difference between narrow evolutionary psychology and the approach described here.

A final difference is that while narrow evolutionary psychology stresses the role of evolution in shaping the human mind and downplays the role of individual learning, ALLNs are as much concerned with individual learning as with evolution. In fact, a crucial objective of many ALLN simulations is to study how evolution and learning interact in creating the adult phenotype. More specifically, since an important biologically evolved human trait is that most human behaviors are culturally learned, i.e., learned from others (Tomasello, 1999), ALLNs are as much concerned with cultural evolution as with biological evolution. Behaviors learned from others are transmitted from one generation to the next with selective reproduction of some behaviors and not of others and the addition of new variants due to learning and invention, imports from other groups, and random errors of transmission analogous to genetic mutations. In this case too, the selective reproduction of behaviors (and ideas, knowledge, value, artifacts) learned from others and the constant addition of new variants result in evolution, but it is cultural evolution rather than biological evolution. While in biological evolution “designs reproduce themselves through the reproduction of the individuals embodying them” (Cosmides & Tooby, 1992), in cultural evolution designs reproduce themselves in large measure independently from the reproduction of the individuals embodying them. Since humans are part of the living world, cultural transmission and evolution are among the phenomena exhibited by the living world and therefore they can be and are addressed using Artificial Life simulations. One important goal of the simulations is to study how biological evolution has created the pre-conditions for learning from others and for other human abilities that make cultural transmission in humans possible and how biological and cultural evolution interact in determining human behavior.

3 What is genetically inherited reflects adaptation to a particular environment but also chance and past history

In Artificial Life simulations, populations of artificial organisms living in a particular environment inherit genetically transmitted information that reflects the properties of the environment and the task(s) the organisms must accomplish in order to survive and reproduce in the environment. The genetically inherited information can code for the morphology of the organisms’ body as well as for their neural network and behavior. For example, if the coloration of the body of a population of artificial insects is encoded in the genotype and the insects are subject to predation from artificial

birds that are more likely to capture a particular individual if the body coloration of the individual is very different from the color of the background environment, the simulation will show an evolutionary change in body color in the population of insects with a progressive assimilation of their body color to the color of the environment. If the environment's color suddenly changes - as actually happened in England at the end of XIX century as a consequence of air pollution due to industrialization - individuals that were well-adapted to the previous environment become ill-adapted to the new environment and a new process of evolutionary change and adaptation begins (unpublished simulations with Andrea Di Ferdinando). (For other simulations of the evolution of both body morphology and behavior, cf. Sims, 1994)

But what is genetically inherited is not only an adaptation to the particular environment in which a species has evolved and reflects not only the properties of this evolutionary environment but it may also reflect chance and the particular past history of the species. This is shown by the following simulations (Lund & Parisi, 1996).

A population of artificial organisms that lives in an environment containing different types of food may evolve a generalist strategy of eating all food types if the energy value of the different food types is the same for the organisms and is fixed. However, if the energy value of the various food types co-evolves with the organisms and therefore it can change, the organisms tend to evolve a specialist strategy of eating only one type of food while ignoring the others. The genotype codes not only for the connection weights of the organisms' neural network but also for the manner in which the organisms process the different food types and extract energy from each of them. Hence, the quantity of energy the organisms extract from each food type can change, and can change in different ways for the various food types, during the course of evolution. In these circumstances evolution causes the emergence (co-evolution) of a better food-processing capacity for one of the food types together with a specialist strategy to eat that type of food in preference to the others.

But consider now what happens if the environment changes and the preferred food type gradually disappears from the environment. While the preferred food is still available, even if in reduced quantities, the population sticks to its specialist strategy. When the preferred food completely disappears while the other nonpreferred foods are still available, there is a sudden shift to a new specialist strategy of eating one of the previously nonpreferred foods. The genetically inherited food-processing capacity of the organisms changes in parallel, with the evolutionary appearance of a better food-processing capacity for the new preferred food.

But in addition to adaptation to the environment and the co-evolution of organisms' traits, it is chance that has a role in determining what emerges evolutionarily and is genetically inherited. One way in which chance enters into Artificial Life simulations is through the initial conditions of the simulation. Evolutionary simulations start with an initial population in which genotypes are randomly generated (from a 'seed') and different replications ('seeds') are used to obtain reliable results. Hence, chance can influence the course of evolution in any individual replication of the simulation and its eventual results. In the simulations just described the initial 'seed' can determine whether a specialist or a generalist strategy emerges evolutionarily when the capacity to process the different food types also evolves, even if with most 'seeds' it is a specialist strategy that emerges. But chance can influence the course of evolution in other ways. In one replication, chance can lead to a preference for one type of food while in another replication another type of food can emerge as the preferred one. Furthermore, when the environment changes and the preferred food gradually disappears, in one replication the initial preference for one type of food can persist until the preferred food completely disappears from the environment whereas in another replication a new preference for a different type of food starts to emerge as soon as the initially preferred food begins to become scarce.

The simulations we are describing also demonstrate the role of previous history and of pre-adaptation in evolution. The way in which a specialist strategy emerges in a population of organisms living in an environment with different food types varies as a function of the preceding history of the population. If we compare the emergence of a specialist strategy for some food type in a population with no previous history and in a population which already has evolved a specialist strategy for another food type and then has abandoned the former strategy because of a change in the environment, we see that the two populations are different in at least two ways. First, the emergence of the new strategy tends to be gradual in the former population and abrupt in the latter population. Second, the population with a previous evolutionary history outperforms the population with no previous history in total energy extracted from the environment. Both differences suggest an important role of pre-adaptation in evolution in that a population with a previous history of alimentary specialism appears to be pre-adapted to evolve a new kind of specialism. A detailed analysis of the activations of single units in the neural networks of the organisms of the pre-adapted population shows how this pre-adaptation is concretely realized in the genetic material and in the resulting phenotype (brain). (For details of this analysis, cf. Lund & Parisi, 1996. For another

simulation in which evolutionary past history explains present adaptation, cf. Phelps & Ryan, in press.)

4 Evolution and learning cooperate to produce adaptation

Imagine a simulated organism that must explore as much of the environment as possible in order to find a target area which cannot be perceived by the organism from a distance (Nolfi & Parisi, 1997). The more of the environment the organism is able to explore the more likely the organism will end up in the target area. The organism is a physical robot with a cylindrical body, infrared sensors distributed on the anterior portion of the body, and two wheels that allow the robot to move around in the environment. The environment is rectangular and it has an outside wall that the robot can sense with its infrared sensors. There is a population of such robots that reproduce differentially as a function of two factors: (1) their ability to find the target area, and (2) their ability to avoid hitting the outside wall. The population evolves in an environment subject to variation between two states: light and darkness. In the lighted environment the robot can sense the wall with its infrared sensors from a greater distance than in the dark environment. The environment changes periodically with a temporal rhythm which coincides with the replacement of a generation by the following generation so that even generations of the population live in a lighted environment and odd generations in a dark environment. Although the organisms do not know if they are living in a lighted or in a dark environment since all they know is how distant they are from the wall at any given time, they behave differently in the two environments. In the dark environment, the organisms tend to include the space near the outside wall in their explorations because it is only when they are very near the outside wall that they can sense the wall. On the other hand, in the lighted environment the organisms are unable to explore the space near the wall because they sense the wall at a greater distance and they don't risk hitting the wall by approaching the wall further (Figure 1, top).

Figure 1. Behavior of two typical individuals of the population without learning (top) and of the population with learning (bottom) in the dark (left) and lighted (right) environment. (From Nolfi & Parisi, 1997.)

These results are obtained when the behavior (i.e., the connection weights) is entirely genetically inherited. If learning during life is added so that an individual can discover in which environment it is living, a different picture emerges. The connection weights which are genetically inherited by each individual are changed during an individual's life as a function of the experience of the individual in the particular environment. (Of course, these learned changes are not themselves

inherited). Since the two environments cause different experiences in the organisms, the organisms now learn to distinguish between them. With learning the organisms' behavior changes. The organisms are less likely to explore the central portion of the environment but they are able now to explore the environment near the wall even in the lighted environment because they now know that, when they first sense the wall in the lighted environment, the wall still lies at a distance (Figure 1, bottom).

This shows one way in which evolution and learning can cooperate to produce a different adaptation to the environment from that which results from evolution alone. In the simulations just described evolution and learning both operate on the connection weights of the neural networks. Evolution selects the initial genetically inherited weights, which already incorporate some useful knowledge about the environment, and learning refines those weights to add some new, more specific, knowledge about the specific environment in which the individual organism happens to live. (For other ways in which evolution and learning can cooperate to find the best connection weights in neural networks, see Nolfi, Elman, & Parisi, 1994; Belew & Mitchell, 1996; Floreano & Urzelai, 2000).

However, the behavior of an organism controlled by a neural network depends not only on the connection weights of the organism's neural network but also on the network's architecture, i.e., the particular pattern of connections specifying which unit is connected with which unit. Although evolution can produce genetically inherited constraints on the learned connection weights and learning can change not only the connection weights but also the architecture, one can advance the general hypothesis that evolution has a more important role in finding the best architecture for a particular type of organism in a particular environment while learning plays a critical role in finding the best connection weights for the genetically inherited architecture (Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett, 1996). In the next section, which is dedicated to another important issue dividing cognitivism and connectionism, the issue of modularity, we will describe a simulation in which evolution creates the network architecture and learning finds the appropriate connection weights for the genetically inherited architecture.

5 Modularity of the mind, modularity of the brain

Narrow evolutionary psychologists tend to think that what is genetically inherited by humans is not just a nervous system and a general tendency/ability to learn but a set of specialized tendencies and abilities, called modules, somehow incorporated in the inherited nervous system. Connectionists, on

the contrary, stress the role of learning and seem to think that not much more than a general tendency to learn from experience is genetically inherited. Most “classical” connectionist simulations are concerned with learning rather than with evolution and do not incorporate specific hypotheses about what is genetically inherited. Hence, connectionists tend to be described as anti-modularist.

ALLNs are different from “classical” connectionism in that ALLNs are not anti-modularist. However, the cognitivist perspective of narrow evolutionary psychology has a different conception of modules than the neural perspective of ALNNs. For cognitivists modules tend to be components of the theories in terms of which empirical phenomena are interpreted and explained. A theory or model of some particular phenomenon can hypothesize the existence of different modules with different structure and/or function which by working together explain the phenomenon. Therefore, cognitive modules are postulated rather than observed entities. For example, in formal linguistics of the Chomskian variety syntax is considered as an autonomous module of linguistic competence in that empirical linguistic data (the linguistic judgements of native speakers) are interpreted as requiring this assumption. Or, in psycholinguistics, the observed linguistic behavior of adults and children is interpreted as requiring two distinct mental modules, one supporting the ability to produce the past tense of regular English verbs (e.g., worked) and the second one underlying the ability to produce the past tense of irregular verbs (e.g., brought) (Pinker, 1999). A third example of postulated modules is the notion of a module so precisely defined in Fodor’s book Modularity of Mind (Fodor, 1983), one of the foundational books of cognitivism.

Neuroscientists also have a modular conception of the brain, but this conception is based on empirical observations of the anatomy and physiology of the brain rather than on theory. The brain is obviously divided up into a variety of ‘modules’ such as distinct cortical areas, different subcortical structures, interconnected systems such as the retina-geniculate-visual cortex or the basal ganglia-frontal cortex system. This rich modularity of the brain, both structural (anatomical and cytoarchitectonic) and functional (physiological), is evidenced by direct (instrumental) observation, by data on localization of lesions in various behavioral/mental pathologies and on neuropsychological dissociations, and more recently, and increasingly, by neuroimaging data.

Neural networks are theoretical models which, unlike cognitivist models, are directly inspired by the physical structure of the brain. In fact, neural networks are at the same time models of the brain and models of the mind. However, the neural networks used in most simulations described in the

connectionist literature are nonmodular. They are homogeneous networks of units with minimal structure usually constituted by an input module (i.e., set of units), an output module, and (almost always) a single internal module in between. This is a serious limitation of current neural network models. The human brain's brilliant performances appear to be due not only to its being made up of one hundred billions of neurons (compared to the tens or hundreds of units of a neural network in a typical connectionist simulation) but also to its being not a homogeneous network of neurons but a very intricately interconnected system of distinct 'modules'. If neural networks claim to be inspired by the structure and functioning of brain, it is not clear how this rich brain modularity can remain absent from the neural networks used in the simulations. Nonmodular networks can illuminate important aspects of mind and behavior because they capture some of the basic physical properties of the brain as a network of units (neurons). However, it is almost inevitable that other important aspects of mind and behavior can only be accounted for if the simulations use modular neural networks that at least begin to match the rich modular structure of the brain.

However, modular networks should reflect the structure of the brain, not the "box-and-arrows" organization of cognitivist models. A module in a modular neural network is a (simulated) physical module, not a postulated theoretical construct. A neural module could be a sub-set of network units with more internal connections linking the units of the module among themselves than external connections linking the units inside the module with units outside the module. Or, more functionally, a neural network module could be an observed correlated activity in a sub-set of the network's units resulting from the pattern of connection weights, without 'anatomical' isolation of the sub-set of units. Neural modules in a neural network can be hardwired by the researcher or they can emerge as a result of evolution or development. If the modular structure is hardwired by the researcher, the researcher tends to be inspired by the actual modular structure of the brain rather than by theoretical considerations based on cognitive models. If the modules emerge spontaneously as part of the simulation, the researcher is interested in ascertaining if the emerging modular structure matches the known anatomical or physiological modularity of the brain.

The real contrast between neural network models and cognitive models, then, does not concern modularity per se but rather the nature of modules. Consider the cognitivist hypothesis that English speakers produce or understand the past tense of verbs using two distinct modules, one for regular verbs and the other for irregular verbs (Pinker, 1999). There appears to be some empirical evidence that these two modules may reside in separate portions of the brain. Patients with lesions in the anterior portion of the brain tend to fail to produce regular past tense forms while their ability to

produce irregular past tense forms is preserved. In contrast, patients with lesions in the posterior portion of the brain tend to show the opposite pattern: they find it difficult to produce irregular past tenses but are able to produce regular ones (Ullman, Corkin, Coppola, Hickok, Growdon, Koroshetz, & Pinker, 1997). This may indicate that two distinct neural modules actually underlie past tense production.

This conclusion can be completely acceptable for a connectionist (at least for ALNNs). What separates the cognitive and the neural approaches to the treatment of past tense is the nature of the modules. Cognitivists claim that the regular past tense module is a rule-based module. To produce the past tense of the verb to work the brain applies the rule: “Add the suffix -ed to the verb root”. In contrast, the irregular past tense module is an association-based module containing a finite list of verb roots each with its associated irregular past tense form. To produce the past of the verb to bring the brain just consults this list of associations, finds the appropriate verb root, and produces the associated past tense form. This theoretical interpretation of past tense behavior is rejected by a connectionist for the simple reason that his or her theoretical tools, i.e., neural network models, do not allow for this interpretation. Neural network models are inspired by the brain, and brains are physical systems made up of physical entities and processes in which all that can ever happen is the production of physico-chemical effects by physico-chemical causes. Hence, a neural network cannot appeal to a rule as an explanation of any type of behavior and cognitive ability. It is perfectly possible that distinct neural modules take care of regular and irregular past tense production but both modules would function in the same basic way, i.e., with neurons influencing other neurons through their synapses. This does not rule out the possibility that one can discover differences in the organization and functioning of the two different neural modules and of course one has to explain why the brain has found it useful to have two separate modules for controlling verb past tense behavior rather than just one, and, more generally, why evolved brains are modular rather than nonmodular.

One reason why the brain is modular might be that modularity solves the problem of neural interference. In most connectionist simulations there is only one task that the neural network must learn. Learning in neural networks consists of adjusting the value of the network’s connection weights so that at the end of learning the network has a set of weights that allow the network to respond to each of the task’s inputs with the appropriate output. Real organisms, however, must be able to accomplish not one single task but many different tasks in order to survive and reproduce. If a network has to learn more than one task and some particular connection is involved in more than

one task, the problem of neural interference may arise. One task may require the weight value of that particular connection to be increased while the other task may require the same value to be decreased. Modularity solves this problem because it assigns separate sets of connections (modules) to each task. In modular neural networks connections are “proprietary” so that it never happens that one and the same connection receives contradictory commands to change its weight value.

Rueckl, Cave, & Kosslyn (1989) have shown that neural networks that have to learn to recognize both the identity of a visually perceived object and the object’s location on the retina learn better if they have a modular rather than a nonmodular architecture. A modular architecture corresponds to the division in the brain between a ventral pathway for recognizing the identity of objects (*What*) and a dorsal pathway for recognizing their location (*Where*) (or for preparing a motor response with respect to the object (*How*)) (Ungerleider & Mishkin, 1982; Mishkin & Goodale, 1995). If one asks how this modular organization originates, evolution is likely to be the answer. In fact, if the network architecture varies from one individual to another individual, is genetically inherited, and a population of networks starts with randomly generated architectures encoded in the genotypes, it can be shown that a modular architecture emerges after a certain number of generations in the population (Di Ferdinando, Calabretta, & Parisi, 2001). Furthermore, in the network architecture which emerges at the end of the simulation the neural module for the *What* task is larger (has more units) than the module for the *Where* task. This agrees with the results of Rueckl et al. (1989) who, in their simulations in which the network architecture is hardwired rather than evolved, have found that the best results are obtained when more hidden units are assigned to the *What* task than to the *Where* task because the *What* task is intrinsically more complex than the *Where* task.

These results are obtained if evolution finds the best network architecture but the organisms learn during life to recognize the identity and spatial location of objects. If both the network architecture and the connection weights are encoded in the genotype, there is no learning during life, and evolution must be able to find both the best architecture and the best connection weights. Under these circumstances, the results are much less good. The evolved organisms make errors in recognizing the spatial location and, especially, the identity of perceived objects, which appears to be due to the fact that the evolved architecture, although it is modular, assigns more hidden units to the easier *Where* task than to the more demanding *What* task. The organisms of the earlier generations tend to reproduce on the basis of their ability to locate the objects spatially (*Where* task) while their ability to identify the objects is not very good (*What* task) since recognizing the identity of objects is not necessary to have a competitive advantage in these early evolutionary stages. When

later in evolution, competition becomes harsher and it would be useful to be able to recognize both the spatial location and the identity of perceived objects, evolution is unable to change the network architecture and to shift hidden units from the *Where* to the *What* task to reach higher overall performance levels.

The *What and Where* organisms must recognize the identity and the location of a perceived object at the same time. A real example is an animal which must be able both to approach prey and to escape from predators: in either case both the identity of the 'object' and its location must be recognized by the animal if the animal is to survive. In other simulations organisms have to accomplish two tasks to survive but in addition they have to be able to choose which of the two tasks to execute at any given time. Cangelosi, Parisi, & Nolfi (1994; see also Cecconi & Parisi, 1993) have simulated a population of organisms that live in an environment containing both food and water and they must be able to find food when they are hungry and to find water when they are thirsty. In any given moment the organisms are informed by their senses about both the location of the nearest food element and the location of the nearest water element and they must decide whether to approach food or water depending on their physiological state. The neural network which controls the behavior of the organisms is genetically inherited and the organisms eventually evolve a neural architecture which includes two modules, a food module and a water module, plus a third module encoding the motivational state of the organism and connected with both the food and the water modules. The motivational state (i.e., the activation pattern of the units constituting the motivational module) can be either hunger or thirst and it shifts to the alternative state after the organism has eaten a certain number of food elements or drunk a certain number of water elements. Based on its current state the motivational module gives control of the network's output to either the food module or the water module. Hence the evolved organisms can be seen on the computer screen to ignore water and approach food when they are hungry and to ignore food and approach water when they are thirsty.

How is the motivational module able to do this? In this architecture the food and the water modules are not entirely anatomical, as they were in the preceding simulation of the *What and Where* task, but they are at least partially functional. As can be seen from Figure 2, the food module and the water module share some of their units and connections. One can distinguish between a food module and a water module because when the motivational state is hunger only some of the network's units have an activation state which covaries with the sensory input (which encodes the location of both food and water) and these units constitute the food module, whereas when the

motivational state is thirst these units are not sensitive to the input and it is the remaining units whose activation state covaries with the input. This is the water module. The motivational module, with its two different activation states encoding hunger or thirst, decides which portion of the sensory input has control of the network's output.

Figure 2. Modular organization of a neural network with distinct modules indicated by thicker connections. (a) Food module. (b) Water module. (From Cangelosi, Parisi, & Nolfi, 1994.)

Another interesting result is evidenced by the evolved network architecture of Figure 2. If one looks at the evolutionary history that has eventually resulted in this network architecture one discovers that this particular population of organisms first evolves an ability to find food when hungry without a comparable capacity to find water when thirsty and only the later generations also evolve an ability to find water when thirsty. Since food and water are equally important for the survival of these organisms, this shows that what is genetically inherited does not reflect only adaptation but also historical contingency. Furthermore, this past evolutionary history, with its intrinsic contingency, has left a trace in the evolved network architecture. The two modules share some network units (unlike cognitive modules, neural modules are not necessarily neatly separated) but the food module which underlies an ability which has emerged earlier is smaller and simpler than the water module which underlies an ability which has emerged later and after the food module was already there. Since the food and water modules share some of their units, it seems as if the water module, when it emerges, co-opts some portion of the already existing food module, which can be interpreted as an example of ex-aptation (Gould & Vrba, 1982).

6 Evolving social behaviors

Many speculations of evolutionary psychology regarding the evolved bases of human behavior concern social behaviors, that is, behaviors elicited by conspecifics or having an effect on conspecifics. In Artificial Life simulations the environment can include nonbiotic elements and organisms of other species but it can also include conspecifics and artifacts created by conspecifics. Therefore, ALLNs can be used to study the inherited architecture of the human mind as this inherited architecture is expressed in social behaviors.

Consider the theory of kin selection (Hamilton, 1964) which assumes that evolution causes altruistic behaviors (behaviors that damage their author and benefit other individuals) to be infrequent if these behaviors benefit a nongenetically related individual, while these behaviors may be routinely

exhibited by an individual if the benefit accrues to a genetically related individual. This can be simulated by having artificial organisms decide to give some of their resources to a conspecific as a function of whether the conspecific is genetically related or not. Sister neural networks in nonsexually reproducing populations are networks which are the offspring of the same parent and therefore they have the same connection weights, except for random mutations, and behave in similar ways. During its life an individual network can encounter both sister networks and other networks and the sensory input tells the network if the currently encountered individual is one of its sisters or an extraneous network. The network's output encodes a decision to give some of the network's resources to the other individual or to refrain from doing so. The results of the simulation show that the behavior predicted by inclusive fitness theory emerges evolutionarily. Evolved individuals give to sisters but not to nongenetically related individuals. By so doing they reduce their individual chances of surviving and reproducing but they increase the survival and reproductive chances of other individuals which have their same altruistic genes. Therefore, altruistic genes and altruistic behaviors directed to genetically related individuals do not disappear from the population.

This approach can be used to study the evolutionary emergence of parental love and filial love. Parental love is defined here as giving one's resources to one's offspring, filial love as the tendency to remain in proximity to parents in order to be able to receive resources from them. Parisi, Cecconi & Cerini (1995) have simulated a population of organisms in which sexually mature individuals (adults) and sexually immature individuals (children) live together. Adults and children exhibit different behaviors. Adults collect the food which is found in the environment while children follow their parents. Newborn individuals live for a fixed length of time as children and then they become adults. Each individual inherits two distinct neural networks, one controlling its behavior as a child and the other one controlling its behavior as an adult. The neural network controlling the behavior of the individual as an adult encodes the location of the single nearest food element as input and responds with a displacement of the individual in the environment as output. When the adult reaches a food element, the food element is captured. The individual's genotype encodes not only the connection weights of its neural network but also a 'parental care' gene which encodes the probability that the individual will eat the food element or give the food element to its offspring. Both the connection weights and the gene's value are inherited with random mutations that can change the connection weights or the gene value. The behavior of the individual as a child is controlled by a neural network with input units encoding the current location of the child's parent (the parent moves in the environment looking for food) and output units encoding a displacement of

the child in the environment.

The behaviors that emerge evolutionarily are the following. Adults exhibit an ability to approach and find food in the environment. They eat part of the food they are able to find in the environment and give the remaining part to their offspring - which is the only way for a child to survive and reach maturity. This is parental love. Since food can only be received by a child if the child is sufficiently spatially close to its parent, children evolve an ability to follow their parents while their parents wander in the environment looking for food; therefore, they tend to remain in proximity to their parents. This is filial love or attachment (Bowlby, 1969). The simulation can be used to identify, given different environmental and other conditions, how much of their food adults tend to give to their already born offspring and how much they prefer to eat in order to survive and generate more offspring.

In these simulations age at maturity is fixed (hardwired). In other simulations one can study how this important developmental parameter can itself evolve and what are the causes and consequences of choosing an early or later age to cease to be a child and become an adult (Cecconi and Parisi, in preparation). In these simulations age at maturity varies from one individual to another individual, is encoded in the genotype and is genetically inherited from one's parents with mutations, and it therefore can change across generations. The results of the simulations show that, if nothing useful for the individual happens during childhood, age at maturity rapidly decreases evolutionarily (it almost goes to zero) so that childhood disappears as a developmental stage and the organisms are already adults when they are born. This results from the fact that only adults reproduce and therefore including a useless nonreproductive stage to one's life reduces the length of adult reproductive life and therefore the chances that one's genes will be present in the next generation.

However, if something useful for the individual happens during childhood, age at maturity stabilizes at some appropriate value and childhood continues to exist. In this variant of the simulation an immature individual learns some skill, say, some food processing ability, which will be used by the individual when it becomes an adult to increase the quantity of energy extracted from the food found in the environment. In this new condition the probability of an individual reproducing will depend not only on the length of its adult life (and, of course, on its ability to find food as an adult) but also on the length of its childhood - since a longer childhood will mean more time to learn to process food and therefore a better food processing ability to be used in adulthood. Again, the simulation can be applied to determine how the age at maturity which stabilizes after a

while depends on various conditions which one can manipulate in the simulation.

But another interesting result of this simulation is the complex pattern of consequences that derive from having evolution determine age at maturity. One consequence is that population size is greater in the population with evolved age at maturity and a useful childhood compared with a population without childhood. Individuals that learn to process food when they are children extract more energy from the food they are able to capture in the environment when they become adults and this allows them to survive longer and have more offspring. (One offspring is generated each time the quantity of energy possessed by a given individual exceeds a given threshold. The individual gives half of its energy to the new offspring.) This causes an increase in population size.

Another consequence is that average life span is greater in the population with childhood than in the population without childhood. (An individual dies either because it has consumed all its energy or because it has reached a maximum age.) Notice that the population with childhood has a greater adult life span than the population without childhood, that is, its adult life span is longer, the addition of a childhood stage to the entire lifespan notwithstanding.

A third, perhaps surprising, consequence is that the individuals of the population with childhood are able to find fewer food elements when they become adults than the individuals of the population without childhood. This is not due to the fact that they are intrinsically less able to find food than the individuals of the population without childhood. When tested in identical (laboratory) conditions, the individuals of both populations exhibit similar levels of genetically inherited food finding ability. The individuals of the population with childhood are able to find fewer food elements because, as we have said, population size is greater in this population than in the population without childhood. Since the environment is the same for both populations, the environment is more crowded for the population with an evolved childhood and therefore each particular individual is able to capture fewer food elements than an individual living in the less crowded environment of the population without an evolved childhood. However, as we have seen, even if fewer food elements are found in the environment, the total quantity of energy extracted from food is greater because of the food processing ability learned during childhood.

The better global adaptation to the environment of the population with an evolved childhood compared to the population without childhood is demonstrated by a final set of simulations in which the two populations live in the same environment and compete for the same resources (food). In 10

out of 10 replications of the simulation in which the initial size of the population with childhood is the same as that of the population without childhood (500 individuals for each population), the population with childhood drives the population without childhood to extinction and it remains as the only surviving population. In another simulation in which the initial size of the population with childhood is half the size of the population without childhood (250 individuals vs 500 individuals), and therefore the former population is at a disadvantage compared to the latter, still the population with childhood outcompetes the population without childhood in 7 out of 10 replications of the simulation.

7 Interaction between biological and cultural evolution

Most human behavior is not only learned but it is learned from others (Tomasello, 1999). Since behavior learned from others is transmitted from one generation to the next and transmission is selective and accompanied by the constant addition of new variants of behavior, this results in cultural evolution, i.e., changes in behavior learned from others in a succession of generations of individuals. ALNN simulations are used to study cultural evolution as much as biological evolution, and how biological and cultural evolution may cooperate in determining human behavior.

A population of individuals lives in an environment that contains a certain number of resource elements. The behavior of each individual is controlled by a neural network with input units encoding the location of the nearest resource element and the output units encoding motor behavior that allows the individual to move in the environment. When an individual reaches a resource element, it captures the resource element and its fitness is augmented by one unit. At the beginning of the simulation the individuals' neural networks are assigned random connection weights. Therefore, the behavior of the individuals of the initial generation is rather inefficient from the point of view of resource procurement. These individuals tend to move randomly or stereotypically and, therefore, they are able to capture a very limited number of resource elements. However, for purely random reasons some individuals happen to have better connection weights than other individuals and these connection weights allow them to behave more efficiently, that is, to capture more resource elements than their less lucky conspecifics.

Just before the individuals of the first generation die (length of life is identical for all individuals and all individuals die at the same time), a second generation of individuals is created. These individuals are born with random connection weights like the individuals of the first generation but,

at the beginning of their life, they are given a chance to learn to procure the resources by imitating the behavior of the individuals of the first generation (Figure 3).

Figure 3. A simple model of learning by imitating another individual. Both the “learner” neural network (left) and the “model” neural network are exposed to the same input and each responds by generating some output. The “learner” neural network learns by using the output of the “model” network (right) as teaching input in the backpropagation procedure. The teaching input tells the “learner” network how to modify its connection weights in such a way that after a certain number of learning cycles the “learner” network responds to the input in the same way as the “model” network.

The best individuals of the first generation are selected as “models” of the individuals of the second generation and, furthermore, when an individual of the second generation learns the behavior of resource procurement by imitating its “model”, some noise is added to the transmission process by slightly (randomly) modifying the teaching input. This noise has the consequence that in most cases a “learner” turns out to be less good at procuring the resources than its “model” but in a few cases a “learner” will outperform its “model”. Of course, the individuals that happen to perform better than their “models” are more likely to be selected at the end of their life as the “models” of the individuals of the next generation.

The process is repeated for a certain number of generations. The capacity to procure the resources tends to increase with each successive generation and after a certain number of generations it reaches a steady state level which is clearly much better than the initial level (Figure 2) (Denaro & Parisi, 1996).

In the simulation just described it is the researcher who selects the best individuals of each generation as the “models” of the next generation. In another simulation the individuals themselves evolve an ability to identify the best individuals of the preceding generation as “models” to be imitated. This ability to select the best individuals of the preceding generation as “models” is genetically inherited. The ability is encoded in the genotype of each individual as a number that measures the individual’s level of ability in selecting good “models”. An individual transmits this ability to its offspring. Biological reproduction is selective in that the individuals that are better at procuring resources (an ability culturally learned from “models”) are more likely to have offspring than less able individuals. When an individual generates an offspring, the offspring inherits the same level of ability to select good “models” of its single parent (reproduction is nonsexual), with the addition of random genetic mutations which can either slightly increase or decrease the offspring’s level of ability to select good “models” compared with its parent’s level.

At the beginning of the simulation all individuals are assigned a random genotype, which means that the ability to select good “models” is not well developed in the initial population. Therefore, in the early generations not much is learned from the “models” since the “models” are not selected appropriately. However, because of the selective biological reproduction and the constant addition of random genetic mutations to the inherited genotypes the average ability to select good “models” progressively increases and at the end of the simulation it is much more developed than it was in the initial population.

The simulation just described is an example of cooperation between cultural evolution and biological evolution. The ability of resource procurement is culturally transmitted (i.e., learned from “models”) whereas the ability to select good “models” is genetically transmitted (i.e., encoded in the inherited genetic material). The two types of evolution are in a reciprocal causal relation. The biological evolution of the ability to select good “models” is made possible by the cultural evolution of the ability of resource procurement. If there were no cultural transmission/evolution of the ability to procure resources, there would be no selective pressure for the evolutionary emergence of the genetically transmitted ability to select good “models”. In fact, individuals that are born with a higher level of the ability to select good “models” are not, by this fact alone, more likely to reproduce. Being a good judge of the quality of potential “models” is of importance only because it allows an individual to learn more, that is, from better “models”. On the other side, if there were no genetic transmission of the ability to select good “models”, the culturally transmitted ability to procure resources would not emerge evolutionarily because individuals would not learn much from “models” because they would be unable to identify good “models”. Hence, biological evolution causes cultural evolution or at least makes it possible.

Another example of an interaction (reciprocal causality) between cultural and biological evolution is the following one. Imagine that to learn by imitating another individual (the “model”) it is necessary for the “learner” to be physically near the “model” - in order to be able to observe the behavior to be imitated. In this simulation each individual has two distinct neural networks. The first network is the network that allows the individual to find the resource elements present in the environment. As in the preceding simulations the connection weights of this network are randomly assigned at birth and the weights are gradually modified as the individual learns to procure the resource elements by imitating a “model”. The second network also allows the individual to move in the environment but the input units of this second network encode the location of “models” rather

then the location of resource elements. The connection weights of the second network are encoded in the inherited genotype and are not changed during an individual's lifetime. As in the previous simulation the individuals that collect more resources are more likely to have offspring and random mutations are added to the connection weights of the second neural network which are encoded in the genotype these individuals transmit to their offspring.

At the beginning of the simulation the connection weights of the second network are randomly assigned to all individuals. Hence, initially, the individuals are unable to move in the environment in such a way that they are able to remain in proximity to "models". The members of the initial generations tend to wander in the environment and to be removed from potential "models". Since in order to learn they must be physically near a "model", this implies that they cannot learn by imitating "models" and therefore they are not very good at procuring the resources. However, the connection weights of the second network are progressively changed by the selective biological evolutionary process and by the genetic mutations and, after a certain number of generations, the members of the population can be seen on the computer's screen to have a tendency to approach and remain in proximity to "models".

In this simulation, too, there is cooperation between cultural and biological evolution. The ability to approach and remain in proximity to a "model" does evolve genetically but not because approaching and remaining in proximity to "models" by themselves increase an individual's reproductive chances. The ability evolves genetically because of the pressure of cultural evolution. Approaching and remaining in proximity to a "model" makes it possible for an individual to learn the capacity to procure resources by imitating the "model". Hence, genetic evolution is dependent on cultural evolution. However, cultural evolution is dependent on genetic evolution in that the culturally transmitted ability to procure resources would not evolve if individuals would not possess a genetically transmitted and evolved ability to approach and remain in proximity to "models" (Figure 4).

Figure 4. Spatial distribution of individuals at the beginning of the simulation and after a certain number of generations. Individuals tend to aggregate in order to learn from "models".

8 Summary

We have described an approach to the study of genetically inherited constraints on human behavior and of the past evolutionary pressures that have resulted in those constraints which differs from the approach of evolutionary psychology in two main respects: the use of neural networks rather than cognitive models for interpreting and explaining human behavior and the adoption of computer simulation as a new way of expressing and testing theories on evolutionary scenarios. In this approach neural networks are viewed as part of Artificial Life, which means that, unlike “classical” neural networks, they have a body, live in an environment, have an inherited genotype, and are members of evolving populations of networks. The approach implies a number of different emphases with respect to evolutionary psychology: (1) evolution is considered as only partially adaptive and the role of nonadaptive mechanisms in evolution emerges from the simulations; (2) the approach is modular in the sense that evolution creates specific adaptations and not only a general learning capacity but modules are physical (structural or functional) components of the brain rather than postulations of cognitive models, (3) learning, and not only evolution, is studied in the simulations, and special consideration is given to the various ways in which learning and evolution can cooperate to make behavior adaptive; (4) the past is important for humans - not only the biological past but also the cultural past; hence, what is simulated is not only biological but also cultural evolution, and their interactions.

The purpose of this chapter can be described as methodological rather than substantial. We have exemplified the approach by describing a number of simulations but these simulations in many cases address generic ‘organisms’ in generic ‘environments’ rather than specifically human organisms living in the human evolutionary environment. To make a more substantial contribution to the study of the specifically human adaptive pattern the simulations must address specifically human behaviors and capacities and they must recreate the specific human evolutionary environment.

References

- Barkow, J.H., Cosmides, L. & Tooby, J. *The Adapted Mind. Evolutionary Psychology and the Generation of Culture*. New York, Oxford University Press, 1992.
- Belew, R.K. & Mitchell, M. (eds) *Adaptive Individuals in Evolving Populations*. Reading, Mass., Addison-Wesley, 1996.
- Bowlby, J. *Attachment*. New York, Basic Books, 1969.
- Buss, D.M. *Evolutionary Psychology. The New Science of Mind*. Boston, Allyn and Bacon, 1999.

- Cangelosi, A., Parisi, D., & Nolfi, S. Cell division and migration in a 'genotype' for neural networks. *Network*, 1994, 5, 497-515.
- Cecconi, F. & Parisi, D. Neural networks with motivational units. In J-A. Meyer, H.L. Roitblat, & S. Wilson (eds), *From Animals to Animats 2*. Cambridge, Mass., 1993.
- Cecconi, F. & Parisi, D. Learning during reproductive immaturity in evolving populations of networks. In preparation.
- Cosmides, L. & Tooby, J. Cognitive adaptations for social exchange. In Barkow, Cosmides, & Tooby, 1992, pp.163-228.
- Cosmides, L., Tooby, J. & Barkow, J.H. Introduction: evolutionary psychology and conceptual integration. In Barkow, Cosmides, & Tooby, 1992, pp. 3-18.
- Denaro, D. & Parisi, D. Cultural evolution in a population of neural networks. In M. Marinaro & R. Tagliaferri (eds), *Neural Nets. WIRN-96*. New York, Springer, 1996.
- Di Ferdinando, A., Calabretta, R., & Parisi, D. Evolving modular architectures for neural networks. In *Proceedings of the Sixth Neural Computation and Psychology Workshop*. New York, Springer, in press.
- Elman, J.L., Bates, E., Johnson, M.H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. *Rethinking Innateness. A Connectionist Perspective on Development*. Cambridge, Mass., MIT Press, 1996.
- Floreano, D. & Urzelai, J. Evolutionary Robots with On-Line Organization and Behavioral Fitness. *Neural Networks*, 2000, 13, 431-443.
- Fodor, J.A. *Modularity of Mind*. Cambridge, Mass., MIT Press, 1983.
- Gould, J.L. & Vrba, E. Exaptation, a missing term in the science of form. *Paleobiology*, 1982, 8, 4-15.
- Hamilton, W.D. The genetical evolution of social behavior. *Journal of Theoretical Biology*, 1964, 7, 1-52.
- Lund, H.H. & Parisi, D. Preadaptation in populations of neural networks evolving in a changing environment. *Artificial Life*, 1996, 2, 179-197.
- Milner, A.D. & Goodale, M.A. *The Visual Brain in Action*. Oxford, Oxford University Press, 1995.
- Nolfi, S., Elman, J.L. & Parisi, D. Learning and evolution in neural networks. *Adaptive Behavior*, 1994, 3, 5-28.
- Nolfi, S. & Parisi, D. Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, 1997, 5, 75-98.
- Parisi, D., Cecconi, F. & Nolfi, S. Econets: neural networks that learn in an environment. *Network*, 1990, 1, 149-168.

- Parisi, D., Ceconi, F., & Cerini, A. Kin-directed altruism and attachment in an evolving population of neural networks. In N. Gilbert & R. Conte (eds), *Artificial Societies. The Computer Simulation of Social Life*. London, UCL Press, 1995, pp. 238-251.
- Phelps, S.M. & Ryan, M.J. History influences signal recognition: neural network models of tungara frogs. *Proceedings of the Royal Society of London. Series B*, in press.
- Pinker, S. *Words and Rules. The Ingredients of Language*. London, Weidenfeld and Nicolson, 1999.
- Rumelhart, D.E. & McClelland, J.L. *Parallel Distributed Processing: Explorations in the microstructure of Cognition. Volume 1: Foundations*. Cambridge, Mass., MIT Press, 1986.
- Rueckl, J.G., Cave, K.R. & Kosslyn, S.M. Why are “what” and “where” processed by separate cortical visual systems? A computational investigation. *Journal of Cognitive Neuroscience*, 1989, 1, 171-186.
- Sims, K. Evolving 3D morphology and behavior by computation. In R. Brooks & P. Maes (eds), *Artificial Life IV*. Cambridge, Mass., 1994, pp. 28-39.
- Tomasello, M. *The Cultural Origins of Human Cognition*. Harvard, Harvard University Press, 1999.
- Ullman, M., Corkin, S., Coppola, M., Hickok, G., Growdon, J.H., Koroshetz, W.J., & Pinker, S., A neural disassociation within language: evidence that the mental dictionary is part of declarative memory and grammatical rules are processed by the procedural system. *Journal of Cognitive Neuroscience*, 1997, 9, 289-299.
- Ungerleider, L.G. & Mishkin, M. Two cortical visual systems. In D.J. Jingle, M.A. Goodale, & R.J. W. Mansfield (eds), *The Analysis of Visual Behavior*. Cambridge, Mass., MIT Press, 1982.




Figure 1




Figure 2