REVISTA MEXICANA DE ANÁLISIS DE LA CONDUCTA 27, 307-336 NÚMERO 2 (DIC.)
MEXICAN JOURNAL OF BEHAVIOR ANALYSIS NUMBER 2 (DEC.)

A BIOBEHAVIORAL APPROACH TO AN ASPECT OF SOCIAL BEHAVIOR

UNA APROXIMACIÓN BIOCONDUCTUAL A UN ASPECTO DE LA CONDUCTA SOCIAL

JOSÉ E. BURGOS[1]

ABSTRACT

The present paper describes a biobehavioral approach to an aspect of social behavior, namely, learning to respond to the ongoing behavior of another individual. The approach was implemented through computer simulations that involved a combination of a neurocomputational model, a network model, a neurodevelopmental model, and a genetic algorithm. In Phase 1 of the core simulation, ten 50-generation lineages evolved under a Pavlovian procedure with one conditional stimulus (CS1). Each lineage had its own random founder population of 100 genotypes. In Phase 2, ten genotypes were randomly chosen from the last generation of each lineage, to form the founder population for a new lineage. In each generation of this lineage, individuals were randomly selected with a small probability to function as 'senders'. Senders were first trained under the same arrangement as their ancestors. Then, they were given 100 maintenance trials under the same arrangement, during which their output activations in the presence of the CS1 served as a CS2 for the rest of the population, which functioned as 'receivers'. All individuals were selected for high conditional responding to their respective CS. Results showed that selection for responding to the behavior of another network reduced population genetic and phenetic variation and increased mean population fitness across generations.

Keywords: social behavior, biobehavioral approach, computer simulations, Pavlovian conditioning, evolution, artificial neural networks, genetic algorithms.

1. Author's note: Correspondence concerning the present paper should be sent to the author, at Centro de Estudios e Investigaciones en Comportamiento, Av. 12 de Diciembre #204, Col. Chapalita, Zapopan, JaliscO 45030, Mexico. For correspondence from the USA, use the following P.O. Box: PMB 30-189, 1605-B Pacific Rim CT, San Diego, CA, 92154-7517. E-mail: jburgos@cucba.udg.mx. Universidad de Guadalajara

RESUMEN

El presente artículo describe una aproximación bioconductual a un aspecto de la conducta social, a saber, aprender a responder a la conducta de otro individuo. La aproximación fue implementada mediante simulaciones por computadora que envolvieron una combinación de un modelo neurocomputacional, un modelo de redes, un algoritmo de desarrollo, y un algoritmo genético. En la Fase 1 de la simulación principal, 10 linajes de 50 generaciones cada uno evolucionaron bajo un procedimiento pavloviano con un estímulo condicionado (EC1). Cada linaje tuvo su propia población inicial de 100 genotipos. En la Fase 2, 10 genotipos fueron seleccionados aleatoriamente de la última generación de cada linaje, para formar la población inicial de un nuevo linaje. En cada generación de este linaje, individuos fueron seleccionados con una probabilidad de .05 para funcionar como 'emisores'. Los emisores fueron primero entrenados bajo el mismo arreglo que sus ancestros. Luego, se les entregaron 100 ensayos de mantenimiento bajo el mismo arreglo, durante los cuales sus activaciones de salida en presencia de CS1 sirvieron de EC2 al resto los miembros de la población, los cuales funcionaron como 'receptores'. Todos los individuos fueron seleccionados por mostrar altas respuestas ante sus respectivos ECs. Los resultados mostraron que aprender a responder a la conducta de otro individuo redujo la variación genética y fenética, e incrementó el éxito reproductivo individual a lo largo de las generaciones.

Palabras clave: conducta social, aproximación bioconductual, simulaciones por computadora, condicionamiento pavloviano, evolución, redes neurales artificiales, algoritmos genéticos.

Many discussions about social behavior revolve around two interrelated issues. First, there is the issue of reduction, that is, of whether or not the social can be reduced to the individual-psychological, or even the genetic-biological. This issue largely arises from the fact that different approaches to social behavior involve different levels of analysis and organization. Second, there is the issue of exactly what qualifies as 'social behavior'. This issue raises questions such as whether or not nonhuman organisms are capable of behaving socially. The vast majority of biologists, especially after Darwin (1859), but definitely after sociobiology (Wilson, 1965), would answer this question emphatically on the affirmative. Psychologists tend to be more adamant in this respect. Some, (e.g., Ribes, this volume), admit the possibility of a kind of 'protosocial' behavior in certain insect species, such as bees, ants, and termites, but impose a qualitative distinction between human behavior and nonhuman behavior, to the point of reserving the term 'social' for human behavior. Others, (e.g., Galef and Schuster, this volume), have no problem in qualifying as 'social' the kind of behavior they study in animals in the laboratory, and would likely join the biologists in this regard.

Conceptual differences notwithstanding, it is clear that an aspect of social behavior that underlies virtually any use of the term, has to do with responding to the ongoing behavior of another individual. Of course, I do not deny the possibility of a single organism responding to its own behavior and, to that extent, of 'social' behavior in reference to a single individual. Nor do I deny the possibility of talking about social behavior in reference to responding to the behavioral *by-products* of another individual (e.g., when a human behaves in relation to something written by another human). However, these possibilities, legitimate and intriguing as they may be, will not be considered here. Rather, I shall focus on responding (specifically, in *learning* to respond) to the ongoing behavior of another individual, and consider it an aspect that is relatively immune to different definitions of the concept of social behavior, that is present in most instances of what most psychologists and biologists would call 'social behavior'.

My emphasis on such a general aspect of social behavior, however, comes with a methodological twist, which has to do with the use of computer simulations as a research tool. According to a widely-held justification, the primary function of computer simulations is to validate their underlying models, through prediction and postdiction. Under this justification, computer simulations are supposed to be realistic, to mirror reality, if their underlying models are to be accepted as empirically true theories. This, however, is not the only possible reason for using computer simulations in empirical science. At least three more reasons can be offered.

First, computer simulations comprise a potentially useful kind of 'thought' or 'imaginary' experiment. Thought experiments constitute a frequent recourse en empirical science. Brown (1991) has referred to them collectively as "the laboratory of the mind" (p. 1), citing as examples Galileo's imaginary experiment on falling bodies (p. 1), Einstein's chasing of a light beam (p. 15), Schrödinger's cat, among other. On the value of thought experiments in biology, Dawkins (1982) has argued that

> ...to understand the actual, we must contemplate the possible... (p. 2).
>
> Playing with an imaginary world, in order to increase our understanding of the actual world, is the technique of 'thought experiment'... At times, thought experiments are purely imaginative and wildly improbable, but this doesn't matter given the purpose for which they are made. At other times they are informed, to a greater or lesser extent, by facts from the real world (p. 3).
>
> ...scientists are sometimes annoyed by the lack of realism in such forms of reasoning. Thought experiments are not supposed to be realistic. They are supposed to clarify our thinking about reality (p. 4).

Under this kind of justification, computer simulations (and their underlying models) are acceptable even if they are unrealistic, insofar as they are useful for clarifying our concepts, explanations, and theories about the phenomenon of interest, as well as our ways of studying it. In fact, an unrealistic simulation (and model) may well be more illuminating in clarifying our thinking about reality than a realistic one. A key argument underlying this reason, in the context of biology, is that the set of actual biological forms is largely due to accident and historical contingency. Hence, many of the lawful regularities at work in the determination of that set can be found only by exploring the much larger set of *possible* biological forms (e.g., Dennett, 1995; Gould, 1989; Emmeche, 1994), which can be done only through simulations.

Under a second reason, computer simulations constitute our best option whenever the subject matter of interest (and/or its causes) is empirically inaccessible to us, because it (and/or its causes) occurs at spatiotemporal scales that are too small too short or too large too long for us. The evolution of a galaxy is a notable example. If we are interested in the studying the evolution of a galaxy, we cannot do it directly. In this sense, the evolution of a galaxy, in and by itself, does not qualify as a phenomenon, in the philosophically technical sense of the term, for it is something that is not manifest in our experience. Our only option thus is to simulate it and our best option is to run computer simulations, based on some mathematical model, and this is, in fact, the way cosmologists proceed. In this manner, computer simulations provide us with surrogate phenomena, so to speak, which become indispensable if we want to study the subject matter of interest at all. They may also be our best option if such a matter is experimentally inaccessible, in spite of it being empirically accessible. In this case, experimental inaccessibility is typically due to the fact that the phenomenon of interest results from complex interactions at different levels of organization, making experimental analyses extremely difficult, and, in many cases, ethically questionable.

A third reason for doing computer simulations, related to the previous one, is that they provide us with an ideal tool for exploring the dynamics of systems of models that have resulted from independent experimental analyses, across disciplines or specialties within the same discipline. Experimental science is primarily analytic, which leads to a fragmentary study of reality. More often than not, scientists from different disciplines (and from different specialties within the same discipline) concentrate on particular aspects of what otherwise constitutes a unitary phenomenon. Models thus are usually proposed to account for (or summarize) specific experimental results that refer to different aspects of the phenomenon of interest. The net outcome of this strategy is a set of models whose logical and conceptual relationships need to be explored, if we are to make any attempt at synthesis (i.e., at reconstructing the phenomenon in question from the conceptual and theoretical pieces that have resulted from the more analytic endeavors). Computer simulations are very helpful for exploring such relationships in a systematic, rigorous manner.

The basic premise of the present paper is that social behavior is behavior. Unless the two occurrences of the term 'behavior' have different meanings, this premise is tautological, but it is precisely this what makes it an ideal premise, for its truth is indisputable. As such, social behavior can, at least in principle, be studied through the methods and concepts from any discipline or specialty whose subject matter is the behavior of organisms. However, behavior has also been studied in a piecemeal manner. So, no discipline can honestly claim to have behavior in general as its subject matter. Rather, behavior is a realm of study that is shared by several disciplines, especially biology and psychology. Behavior, then, is a subject matter of biology as much as of psychology. In all fairness, practitioners of these disciplines can only claim to have a certain aspect of behavior as its subject matter. Particularly, biologists have typically focused on the underlying biological structures and processes of behavior, as well as on its adaptive value in ancestral environments, while psychologists have concentrated their efforts more on the causes of behavior in present environments. Some practitioners of both disciplines have found common grounds in a discipline called 'behavioral neuroscience', which has produced a journal of the same name. Behavior (and, hence, social behavior) is, *mutatis mutandis*, like the proverbial elephant of the oft-mentioned Indian fable, while biologists, psychologists (and, to some extent, sociologists and economists) are like the blind men.

Social behavior is a kind of behavior that seems (as many other kinds of behavior) suitable for study through computer simulations. First, there is considerable conceptual confusion about it and its determinants. Second, it results from complex interactions at different levels of organization, which makes it a particularly difficult phenomenon to be studied experimentally. When viewed from the standpoint of the theory of evolution by selection, some of its causes (as much as some of the causes of any other kind of behavior) are empirically inaccessible to us, insofar as they are ancestral. Third, it is a phenomenon whose study has been divided into different aspects, some of which have been studied by biology, while others have been studied by psychology. On this basis, computer simulations may help us clarify our thinking about social behavior, by providing us with surrogate phenomena that emerge from the dynamics of a system of models that synthesizes analytic efforts from biology and psychology.

In view of the above considerations, the conceptual, theoretical, and methodological intersections among biological and psychological approaches to social behavior pose a formidable unification challenge. Given the complexity of the topic, my purpose in the present paper is very modest. Specifically, the synthetic strategy I shall adopt is to use computer simulations to explore the dynamics of a system of biological and psychological models, in relation to that aspect of social behavior I shall concentrate throughout the present paper (viz., learning to respond to the ongoing behavior of another).

The present approach adopts a simplifying characterization of individual organisms as *artificial neural networks* whose structure is largely heritable (in the

population-genetic sense of the term), and whose behavior can confer an adaptive advantage, in the sense of being *selectable for* and, hence, evolvable, at least in principle. The empirical part of the paper is a set of simulations in which artificial neural networks are selected for learning to respond to the behavior of other artificial neural networks. In the first section, I describe the system of models that underlies the simulations. The simulations are described in the second section. In the third section, I conclude with some implications for social behavior in general.

Like any approach to complex phenomena, the present one is preliminary, tentative, and most incomplete. I just want to provide a starting point for what could be called, for lack of a better term, 'computational sociobiology', the study of the biological bases (phylogenetic as well as ontogenetic) of social behavior in artificial systems. The term 'sociobiology' has acquired a rather unfortunate fame, partly due to misrepresentation, partly due to the way some sociobiologists have developed their field, especially when it comes to applications to certain kinds of human social behavior. None of such misrepresentations or applications, of course, are intended here, so I do not pretend to derive any direct implications for human social behavior. However, as Sober (1993) has pointed out,

> ...there is no 'magic bullet' that shows that sociobiology is and must remain bankrupt, nor any that shows it must succeed. Any discussion of the adequacy of sociobiological models inevitably must take the models one by one and deal with the details. (p. 185).

A SYSTEM OF MODELS

A fundamental guiding idea in the present approach is that an adequate understanding of behavior in general and, hence, of social behavior in particular, requires a synthetic, unifying, integrating effort that should aim at a system of biological and behavioral models. The key notion of a *system of models* refers to a set of interrelated mathematical models, related in such a way that the outputs of (numeric solutions to) some models serve as (numeric) inputs to others. Such a notion arises from considering behavior as a *unitary phenomenon* that has been experimentally studied and theoretically accounted for in a piecemeal, analytic manner. In this sense, models of different aspects of behavior (and, hence, social behavior) constitute different pieces, so to speak, to the puzzle of behavior as a unitary phenomenon.

My intention here, of course, is not to assemble the entire puzzle, but only a small fraction of it, by selecting and shaping certain pieces so that they fit together in a logically sound and empirically fruitful (via computer simulations) manner. Each piece inevitably involves a great deal of simplification, for, after all, the whole idea of a model is to simplify. So, the models do not attempt to capture all (not even a substantial fraction) of the richness of their motivating phenomena. Hence, the resulting fraction of the puzzle will present a rather rough, abstract figure. The

models presented here have been described in more detail somewhere else (Burgos, 1995, 1997; Burgos & Donahoe, in press; Donahoe, Burgos, & Palmer, 1993; Donahoe *at al*, 1982; Donahoe & Palmer, 1994; Donahoe, Palmer, & Burgos, 1997a, 1997b). Here I only give a very brief, informal description.

The basic prescription underlying the system of models I present here is quite simple. In order to understand behavior (social behavior included) as a unitary phenomenon, we must take into account the structure and functioning of nervous systems, the structure and functioning of how the environment controls the behavior of organisms, how genotypes become nervous systems, and how behavior in such environments affects reproductive success. If one believes in the possibility and necessity of unifying biology and psychology (and I must certainly do), then the need to take all of these aspects of behavior into account can hardly be denied. Of course, this is not a new idea. The necessity of synthesis has been expressed by several authors (e.g., Hollis, 1984, 1990; Kamil, 1994; Kandel & Schwartz, 1982; Teitelbaum, 1977). However, the particular forms that these attempts acquire vary widely from one author to the other. To be sure, my own attempt is different from the ones usually found in the biological and psychological literature. As a computational synthesis, my proposal is more readily identified with those found in the very young (but rapidly growing) discipline known as 'Artificial-Life', or 'A-Life' (e.g., Adami, Belew, Kitano, & Taylor, 1998; Brooks & Maes, 1994; Langton, 1992, 1994; Langton & Shimohara, 1997).

Even the most convinced biopsychologists or biobehaviorists, as we could call those who believe in the necessity of said unification, might be somewhat skeptical with respect to the *possibility* of achieving it. In effect, each one of those fields involves a lifetime of research, so any attempt to unifying them is in high risk of becoming a never-ending endeavor. Anyone can certainly spend a lifetime studying only the relationships between environment and behavior, even studying a fraction of these relationships. Therefore, any attempt to unify this field with developmental neuroscience, neuroanatomy, and cellular neurobiology (each one as extensive as the other), and then trying to unify all of these fields would certainly seem impossible to achieve, no matter how strongly we believe they should be unified. And the unification enterprise becomes even more daunting when we take the phylogenetic, evolutionary aspect into account.

Again, the idea is not to attempt a full unification, one that takes into account every single known concept, phenomenon, theory, and phenomenon found in all of these research fields. On the contrary, a synthetic endeavor must be highly selective and simplifying, if it is to accomplish anything at all. Nonselective, all-encompassing unification efforts are doomed to wander hopelessly and aimlessly in the vast ocean of scientific knowledge. I certainly do not pretend to achieve a final complete unification, but only a tentative partial one. As I use the term here, 'unification' does not imply 'completeness'. Rather, it refers to relating certain notions that typically remain unrelated. Also, there is, once again, the issue of reduction. My use of the term 'unification' here does not intend to denote

(nor connote) 'reduction'. Consequently, criticisms to the present synthesis attempt that are based on a rejection to reductionism, simply do not apply. In fact, I will argue that reduction (understood in a certain way) of the behavioral (or psychological) to the neural is not possible in the kind of artificial system presented here.

Another clarification, before I describe the models, is that an emphasis on synthesis does not entail a rejection of analysis, be it experimental, conceptual, or theoretical. On the contrary, synthesis absolutely and inevitably arises from all these forms of analysis. I am a true believer in analysis as the best way of achieving some understanding about almost anything. However, analysis is not the end of the road. I see analysis in science only as a *means* that will be relevant to the extent that it leads to synthesis. After all, the starting point of science was nature as such. The analytic character of science necessarily involves a detour that, more often than not, takes us far away from nature. However, science must always attempt to find its way back to its original motivation, nature as such, and the only way is through synthesis. To this extent, synthetic science will be as legitimate as analytic science. Synthetic science, however, must be carried out wisely, very carefully, systematically, and very patiently. Otherwise, it could well end up farther away from reality than analytic science.

The neurocomputational model

The neurocomputational model describes the structure and functioning of the neural processing element (NPE), a sort of abstract neuron that receives input signals from its local environment, processes such signals, and returns an output signal. Fundamentally, then, the NPE is an input-processing-output device inspired by the basic structure and functioning of neurons.

The neurocomputational model consists of two components, namely, the activation rule and the learning rule (see the Appendix in Donahoe, Burgos, & Palmer, 1993, for a mathematical description). The activation rule determines the NPE's state (represented as a real number between 0 and 1) at a moment in time (the model is a discrete-time one), as a function of its input signals at that moment and the strengths at which the carriers of such signals are connected to the NPE. Such strengths are measured as connection weights (also represented as real numbers between 0 and 1). In the jargon of cellular neurobiology, the NPE is analogous to a neuron's dendrites plus soma, the carriers of its input signals are analogous to presynaptic processes, and a connection weight is analogous to the synaptic efficacy of one of such processes. Input signals to an NPE are partitioned into excitatory and inhibitory. These signals are processed separately by the NPE up to a point, after which they are combined to determine the NPE's activation state. A similar strategy is used in the activation rule of the model known as the 'cognitron' (Fukushima, 1975).

The learning rule determines how connection weights change from moment to moment, as a function of changes in the NPE's activation state and that of its presynaptic inputs. Changes in connection weights also depend on a magnitude we refer to as a *diffuse reinforcement signal*, which intends to simulate the kind of dopaminergic signals that have been found to be critical for the biological effect of primary reinforcers (food, water, etc.). This kind of signal is diffuse in that it is exactly the same signal for all connections at any given moment in time, although it may differ from one moment to the next. The learning rule has two mutually exclusive modalities, namely, incremental, which is enabled when the discrepancy signal is positive, and decremental, which is enabled whenever the reinforcement signal is zero or negative. Finally, the functioning of the activation and learning rules is modulated by certain free parameters, which, from the model's perspective, are considered as properties of the individual NPE.

The network model

A neural network is a set of interconnected NPEs. The *network model* consists in a taxonomy of NPEs plus a number of general prescriptions on how they are to be interconnected. The taxonomy in question arises largely from the position that different NPEs occupy in a network, and it is based on the traditional connectionist distinction between input, hidden, and output layers of elements. Strictly speaking, input elements are not NPEs, for their activation states are not computed using the activation rule, nor do they receive signals from other NPEs. Rather, the input elements represent the sensors or receptors of a neural network, for which their activation represents the occurrence of an environmental sensory stimulus (e.g., a light or a tone). For the present purposes, input elements were subdivided into *s1* and *s2*, which represented different sensory channels or modalities (e.g., vision versus audition). Only hidden and output elements, then, qualify as NPEs.

Hidden elements are divided into cortical and subcortical. Hidden cortical NPEs are subdivided into sensory-association (*sa*) and motor-association (*ma*) NPEs, which are roughly analogous to neurons that constitute the cortical areas of the same name. Hidden subcortical NPEs are subdivided into *ca1* (which are named after the hippocampal area of the same name, and receive plastic connections from *sa* NPEs) and *vta* (which are named after the dopaminergic nucleus known as ventral-tegmental area, and receive plastic connections from the *ma* NPEs). The activation of these NPEs across successive moments in time provides the source of the diffuse reinforcement signals. In addition to receiving plastic connections from the *ma* input elements, *vta* NPEs receive a nonplastic, initially strong connection from a *us* input element, whose activation represents the occurrence of a primary reinforcer.

Functionally, the distinction between *sa* and *ma* NPEs is based on the source of the diffuse reinforcement signal that modulates changes in the weights of their respective presynaptic connections. In the case of *sa* NPEs, the signal source is

the activation of the *ca1* NPEs across successive moments in time. The *ca1* signal is amplified by the *vta* one. This latter signal also modulates changes in the weights of the presynaptic connections to *ma* NPEs. Changes in the weights of the presynaptic connections to *ca1* subcortical NPEs (projecting from the *sa* NPEs) are modulated by the *ca1* signal (amplified by the *vta* signal). Changes in the weights of the presynaptic connections to *vta* NPEs (projecting from the *ma* NPEs) are modulated by the *vta* reinforcement signal.

As a restriction, only *sa* and *ma* NPEs could be excitatory or inhibitory. Also, only inhibitory connections among NPEs within the same layer were allowed, which implemented a sort of lateral inhibition. The typical configuration of lateral inhibition involved at least three NPEs within the same layer, two excitatory and one inhibitory. An excitatory NPE connected to an inhibitory NPE could inhibit another excitatory NPE if the first excited the second and the second was connected to the third NPE.

Finally, output NPEs are subdivided into *R* NPEs, whose activation represents the occurrence of an operant response, and *CR/UR* NPEs, whose activation represented the occurrence of a either a conditional response (CR) or an unconditional response (UR). The distinction between CR and UR depended on whether the activation of a *CR/UR* output NPE was caused by the activation of the *s1* and/or *s2* input elements, or by the activation of the *us* input element, respectively. The only difference between *R* and *CR/UR* NPEs was that the latter received nonplastic, strong connections from the *vta* NPEs, which, in turn, received a direct, nonplastic, strong connection from the *us* input element. The activation of the *us-vta-CR/UR* path, then, simulated the occurrence of an unconditional reflex.

A typical neural network is organized into layers through which input activations are propagated in a feedforward manner from inputs, to hidden, to output NPEs. A network may have only one layer of input elements, one or more layers of *sa* NPEs (each with its own layer of target subcortical *ca1* NPEs), one or more layers of *ma* NPEs (each with its own layer of target subcortical *vta* NPEs), and only one layer of output NPEs. Input elements project (i.e., are connected) to the NPEs that constitute the first *sa* layer. If a network has more than one *sa* layer, then *sa* NPEs in one layer project to the *sa* NPEs that constitute the immediately succeeding layer. The NPEs that constitute the last *sa* layer project to those constituting the first *ma* layer. Here is where the sensory-motor link takes place. If a network has more than one *ma* layer, then NPEs constituting one layer project to the NPEs constituting the immediately succeeding layer of *ma* NPEs. Finally, NPEs constituting the last *ma* layer project to the NPEs constituting the output layer. Sensory-association NPEs also project to *ca1* NPEs, while *ma* NPEs project to *vta* NPEs. The magnitude of the diffuse reinforcement signals thus depends not only on the occurrence of a primary reinforcer (as given by the activation of the *us* input element), but also on the occurrence of sensory, neutral stimuli (as given by the activation of *s1* and *s2* input elements). In this manner, the model implements a sort of internal mechanism for secondary reinforcement. After some learning has

occurred (i.e., after certain environmental conditions have caused changes in the appropriate connection weights), the mechanism in question allows sensory input elements to activate (through the *sa* and *ma* NPEs) the *ca1* and *vta* NPEs, thus producing the diffuse reinforcement signals in the absence of a primary reinforcement.

In the present approach, a neural network represents an individual's structural or relatively fixed phenotype, analogous to the structure of the nervous system of a natural organism. The functioning of this phenotype, as given by the activation patterns of its constituting elements, represents the individual's behavior (or performance or functional or variable phenotype). An individual's behavior at a moment in time thus is formally defined by the vector of activations of *all* of its constituting NPEs at that moment. However, when this definition is applied to actual biological forms, it amounts to defining an organism's behavior at a moment in time as the state of its *entire nervous system* at that moment. Taken literally, such a definition imposes a rather massive (and largely unattainable) observation requirement. When an organism's behavior is observed across time, even in a simplified environment, it is impossible to keep track of the state of its entire nervous system. Even in a small artificial neural network, which is much simpler than the simplest nervous system known to date, such an observation task becomes cumbersome. So, in the context of behavioral studies, the behavior of an organism is typically defined in terms of the *primary motor output* of the organism's nervous system. And this is the strategy I will follow here. More precisely, the behavior or performance of a neural network at a moment in time is defined as the state of its output activation vector at that moment. Also, learning can be conceptualized as a change (either incremental or decremental) in one or more connection weights, as a function of changes in the network's environment and as given by the learning rule. Thus we obtain a connectionist interpretation of the learning-performance distinction.

The environment model

In the present approach, the environment is conceptualized basically in the same manner as it is conceptualized in animal-conditioning experimental research. The environment is thus divided into stimuli, whose occurrence is represented by the activation of a network's input elements. Two general types of stimuli are identified, namely sensory stimuli and primary reinforcers. This distinction is analogous to the one made between CSs and USs in Pavlovian conditioning research, and between discriminative stimuli and primary reinforcers in operant conditioning research. Sensory stimuli are subdivided into modalities, which are roughly analogous to the ones identified in conditioning research (e.g., visual, auditory, etc.). Formally, a stimulus is defined as a vector of sensory-input activations. A response is defined in terms of the activation level of a kind of output element.

The environment is also defined in terms of the statistical and temporal relationships that the different kinds of events (stimuli and responses) can maintain. In this manner, we can define a Pavlovian procedure as an arrangement in which the *us*-input activation is temporally and statistically dependent on the activation of the sensory input elements. An operant procedure can be defined as an arrangement in which the *us*-input activation is temporally and statistically dependent on the activation of the sensory-input elements and on the level of activation of the *R* output elements. As a starting point, I shall focus on Pavlovian arrangements and, hence, on that aspect of a network's behavior that is given by the activation of the *CR/UR* NPEs in the presence of certain sensory stimuli. Also, for my present purposes, I prefer not to adopt any particular interpretations of what input and output signals exactly represent in the real world.

The neurodevelopmental model and the genetic algorithm

Formally speaking, a neural network can be viewed as a system of multiple concurrent realizations of a neurocomputational model, such that the outputs of certain realizations provide the inputs to others. This way of defining a neural network provides us with a rather unique criterion for evaluating a neurocomputational model, a criterion that for all ends and purposes transcends the model itself. Model plausibility (or validity) constitutes a primary concern in modeling work. Usually, plausibility (however we choose to define it) is viewed as an internal property of the model. For example, a model is regarded as being more or less plausible (if we view model plausibility as a matter of degree) when it corresponds more or less to certain data. When multiple realizations of the model are put together, making them interact with one another, we can observe the global behavior of the resulting system. However, such a behavior may depend, at least in principle, not only on the model itself, but also on how many realizations of it the system involves and how they are related. These variables are not part of the model itself, and yet they can affect the behavior of a system composed of multiple realizations of it. When this happens, I call the system's global behavior an *emergent phenomenon.* This possibility raises the issue of whether or not the behavior of such a system can be reduced to the behavior of its components. It also makes us revise our model-plausibility criteria.

With respect to reduction, it is clear that if a system's behavior not only depends on the behavior of its constituting parts, but also on the way such parts are put together, then the former behavior cannot be reduced to the latter, at least not completely. The reason for this is simply that the system's variables (number of parts of how they are interconnected) are not variables in the model that describes the functioning of the parts in isolation. Hence, the system's global behavior cannot be deduced from the part's local behavior, in spite the fact that the former clearly depends on the latter. However, such a dependency is only partial. If we adopt a multicausality scheme, admit the possibility of causal

relationships between a system's behavior and that of its parts, and accept (in a move that revisits Aristotelian causality) the structure of the system as a legitimate cause of its global behavior, then the behavior of the part determines the behavior of the system only partially. In this case, the model that describes the functioning of the part is necessary but not sufficient to understand the system's behavior, even if the system consists of multiple realizations of the same model. Therefore, reduction is not possible, at least not in the sense of a system's behavior being completely deducible from the functioning of its parts (cf., Bickle, 1998).

With respect to plausibility, if the behavior of a system seems implausible (relative to some proposed analogous system of interest), then it is always possible that the source of the implausibility is the way the system is put together, rather than the model that describes its parts. That is to say, a system may behave implausibly, even if the model that describes the functioning of its parts is plausible (relative to some other proposed analogous system). Plausibility thus can breed implausibility. A model that is plausible at its own level of analysis may generate implausible behavior in a system consisting of multiple realizations of the model in question, again, if said behavior also depends on the way such realizations are put together.

Applying the above considerations to artificial neural networks involves asking whether or not a network's behavior depends on the network's structure. As far as I know, only the present neurocomputational and network model has shown clear and systematic dependencies of that kind. Moreover, such dependencies are consistent with certain conditioning phenomena (e.g., the ISI function and variations in the optimal ISI across preparations in Pavlovian conditioning; see Burgos, 1995, 1997; Burgos & Donahoe, in press). The main implication for the present purposes is that a network's architecture is a variable that affects its behavior. This implication leads us to always consider the possibility that when a network behaves implausibly (relative to some natural animal organism), such behavior may be due to the network's architecture. So failure to simulate a given behavioral phenomenon in a neural network may not be due to the neurocomputational model. This possibility raises a formidable methodological problem, for we are forced to search the rather massive architecture space (the set of all possible neural networks), in order to determine the extent to which the network's behavior depends on its architecture.

Clearly, it is most impractical to search such a large space manually. An automatic search procedure thus is needed. Ideally, the procedure in question should not only allow for a relatively efficient and successful search, but also it should be consistent with some biological theory. The ideal choice is a genetic algorithm, a model of evolution by selection, inspired by the synthetic theory of evolution. Before I describe this model, however, I need to mention briefly the model I have used to simulate neural development, a key process where ontogeny and phylogeny intersect in critical ways.

The *neurodevelopmental model* was an algorithm for transforming bit strings (representing genotypes, the individuals' genetic makeup) into neural networks. The use of bit strings to represent genotypes is taken from the theory of genetic algorithms (see later). The bit string used in the present model consisted of 323 bits representing gene loci, whose possible values (0 or 1) represented alleles. Each bit string was subdivided into 68 fragments, each one encoding in binary a decimal parameter that determined some feature of a network's architecture. The three basic features were the maximum numbers of each kind of element (*s1*, *s2*, *sa* excitatory and inhibitory, *ca1*, *ma* excitatory and inhibitory, *vta*, *R*, and *CR/UR*), the probabilities of different kinds of NPE connecting to one another, and the free activation and learning parameters for each kind of NPE. The model was inspired by the general organization of neural development into functional stages, found in developmental neuroscience (e.g., Brown, Hopkins, & Keynes, 1991). According to this organization, neural development starts with the stage of proliferation, during which immature nerve cells (neuroblasts) are produced through mytosis. Then, newborn cells migrate from their point of origin (the so-called 'ventricular zone) to their final resting place (the so-called 'cortical plate'). In a developing neural network, migration amounts to distributing nondifferentiated elements into layers. After migration, the elements are differentiated, which amounts to deciding probabilistically whether an element will be excitatory or inhibitory, or whether an output element will be *r* or *CR/UR*, and so on. The activation and learning free parameters are also assigned to each kind of NPE during this stage. During the synaptogenesis stage, elements of different kinds are probabilistically connected, and the initial weights are computed by multiplying the spontaneous activation of the potential pre-and postsynaptic elements (i.e., a Hebbian rule was used to determine the initial weights). Hence, initial connection weights were not directly encoded in the genotype. Finally, during the cell-death stage, elements that did not receive any connections (except for input elements) and/or elements that did not send any connections (except for subcortical and output NPEs) were eliminated from the network. Because many of the encoded parameters were probabilistic, the same genotype could produce many different networks. So there was a one-many relation between genotypes and structural phenotypes, which allowed for phenotypic variation associated with the same genotype.

Finally, a genetic algorithm was used as a computational procedure for simulating evolution by selection, according to the synthetic theory (see Goldberg, 1989; Holland, 1992; Mitchell, 1996, for introductions to genetic algorithms). The core of a genetic algorithm is the so-called 'selection scheme', a procedure for selecting individuals (whose genotypes are represented as bit strings that encode certain phenotypic traits) for reproduction, based on a quantitative criterion known as 'fitness'. This criterion is used for selecting which individuals (and with what frequency) will reproduce and, hence, pass their genetic material to the next generation. The genotypes from those individuals that have been selected for reproduction are recombined, with a certain, usually high probability, in order to

produce new descendant genotypes, which will constitute members of the next generation. Then, genes in the newly-produced genotypes are mutated (i.e., if a gene with an original allele of 1, then is changed into an allele of 0, or viceversa) with a small probability. In a genetic algorithm, then, most of the genetic variation in a population arises from genetic recombination or crossover at the moment of reproduction.

For simplicity, the present neurodevelopmental and genetic algorithm assume that all networks had a built-in *us-vta-CR/UR* path whose activation simulated the occurrence of an unconditional reflex. Networks thus had the capacity of showing one and only one unconditional reflex. Therefore, the present approach already assumes the occurrence of a previous, hypothetical evolutionary process that resulted in that kind of structure. In this sense, the approach represents an incomplete depiction of the phylogeny of behavior, one that assumes the availability of certain primary anatomical structures that are necessary for learning to occur and that must have been the result of evolution by selection.

In the present approach, an individual's fitness is defined as its behavior under certain environmental conditions. More precisely, individual fitness is identified with the behavior (functional phenotype) of a neural network (structural phenotype developed from a genotype) under certain Pavlovian contingencies. An individual thus is conceptualized as being constituted by a genotype (defined as a bit string), a structural phenotype (defined as a particular neural-network architecture), and a functional phenotype (defined as the activation of the network's output NPEs).

SIMULATIONS

General description

The main aim of the simulations presented here was to apply the system of models that I described in the preceding section, in order to simulate the phylogeny of a certain aspect of social behavior, namely, responding by an individual to the behavior of another individual, under a Pavlovian arrangement. My focus on this particular kind of arrangement is a convenient methodological simplification, more than a theoretical proposal. Indeed, a neural network consists of a set of interconnected NPEs. Although many different network architectures are possible, the functioning of any NPE is described by exactly the same neurocomputational model. Therefore, the present approach does not introduce any fundamental (i.e., principled, theoretical) distinction between Pavlovian and operant conditioning (see Donahoe, Palmer, & Burgos, 1997a, 1997b). That is to say, when this distinction is reconstructed within the context of the present neurocomputational and neural-network models, the only meaningful separation ends up being methodological in nature (Burgos, 1999).

The Pavlovian arrangement used in all of the simulations involved a forward-delay procedure with continuous reinforcement. Such an arrangement involved activating a network's sensory input units with some level of activation for 7 time-steps (ts), where a ts represents a relatively short moment in time. The specific activation of the input units depended on whether the CS in question represented a nonsocial CS (CS1) or a social CS (CS2). CS1 was defined as a constant activation of input units of a given kind (either *s1* or *s2*) with a value of 1, throughout those 7 ts. This CS was considered as nonsocial only in the sense that its source was not the output activation of another neural network. In contrast, CS2 was defined as the activation of a given type of input unit with a level equal to the average activation across *CR/UR* output NPEs of another, randomly-chosen network. In all the simulations, CS1 involved the activation of all *s1* input units, whereas CS2 involved the activation of all *s2* input units. Reinforcement involved the activation of the *us* input unit with the maximum possible value.

Any given network received instances of only one kind of CS, although different networks within a population could receive different kinds of CS. Networks that received CS1 were called 'senders', while those that received CS2 were called 'receivers', both roles being mutually exclusive and permanent within a generation. In all simulations, the probability that a network was chosen as a sender was .05, for an average of 5 senders per generation and 19 receivers randomly assigned to each sender. Having such a small proportion of senders was motivated by the aim of making most of selection (and, hence, most of evolution) depend on the behavior of the receivers, the behavior of interest here. Once a network was chosen as a sender, it remained as such throughout its entire generation. The same applied to the receivers. For all instances of CS1 and CS2, reinforcement was scheduled to occur at ts = 7. Hence, the interstimulus interval (ISI) was 6-ts long. The intertrial interval was not explicitly simulated. Rather, it was assumed to be sufficiently long as to allow all NPE activations in a network to decay to a low (close-to-zero) level of activation.

A simulation amounted to creating a lineage, defined as a sequence of populations related by descent, each one representing a generation. All populations consisted of 100 individuals. All lineages started with a founder population (known as G0) consisting of randomly-generated genotypes. A generation was defined as a population of 100 individuals that underwent a development-training-selection-reproduction cycle. Individual selection depended on the individual's fitness, which was defined for all simulations as the average activation across its *CR/UR* NPEs at ts = 6 (the moment immediately before reinforcement), across 25 additional trials given after the 100 training ones. The selection scheme used for all generations was the so-called 'tournament selection'. This scheme consisted in choosing, randomly and with replacement, samples of 5 different individuals from the population, and selecting for reproduction the individual with the highest fitness. If all individuals had exactly the same fitness, then no selection took place for that tournament, and a new tournament took place, until a reproduction pool

of 100 individuals was obtained. Then, reproduction took place between pairs of individuals randomly chosen from the reproduction pool.

For all generations, the genotypic crossover probability during reproduction was set to .8, while the mutation probability of newly-produced genotypes, after reproduction was .001. Training only modified a network's connection weights. However, an individual's experience did not modify its own genotype, which ensured a non-Lamarckian genetic transmission. Finally, there was no overlap among generations, that is, parent networks did not interact (behaviorally or otherwise) with their offspring (i.e., they were taken away from their parents at birth). After a new population of 100 genotypes was obtained, the development-traning-selection-reproduction cycle started anew, and it was repeated for either 50 or 100 generations.

Simulation 1

In the first simulation, I simply tried to simulate the evolution of responding to another individual's behavior from scratch, so to speak, that is, starting directly with a randomly-generated population of genotypes. A single lineage was generated, consisting of 100 generation of 100 individuals each. Each sender was first given 100 acquisition CS1 (reinforced) trials. This initial training was considered necessary in order to obtain a relatively high *CR/UR* activation at ts = 6, as to serve as a substantial CS2 input signal for the receivers. Typically, naive networks (i.e., networks that have not received any training whatsoever) show very small (close-to-zero) output activations. Therefore, naive senders provide a very weak input signal for the receivers, thus precluding the latter from achieving any substantial learning. This difficulty is overcome by pretraining the senders. The implication of this for social behavior in general is that social learning is far more effective if at least one of the participants has had some previous experience. So, after pretraining the senders, they were given 100 additional (maintenance) CS1 trials, corresponding (in discrete time) to 100 CS2 acquisition trials for the receivers. Finally, all networks were given 25 additional reinforced trials (CS1 for the senders and CS2 for the receivers), in order to compute individual fitness. After selection for reproduction, senders could mate among themselves or with receivers. Also, receivers could mate among themselves. Given the small number of senders in the population, however, the most likely mating relation was the one among receivers, for they comprised 95 percent of the population on the average.

The main result was that there was no substantial increase in the mean population fitness across generations and, hence, there was no substantial evolution. This result was due to the following situation. At the outset of the lineage (i.e., in G0), the vast majority of the individuals showed very little acquisition, for which their *CR/UR* activations tended to be close to zero. Because of this, the senders showed a very little *CR/UR* activation, even after 200 reinforced trials (in preliminary simulations, this situation did not improve with larger numbers of

training trials). Therefore, CS2 was a very weak signal, which made acquisition far more difficult for the receivers. Consequently, most of the population showed very little or no acquisition at all. Of course, there must have been some behavioral variation within G0, however small. Otherwise, selection would not have taken place and the simulation would not have gone beyond G0. Nevertheless, since selection and reproduction occurred mostly among the receivers, there was very little opportunity for any substantial evolution due to any learning achieved by the senders. Any adaptive advantage that learning may have given to the senders was statistically swamped at the moment of selection and reproduction by the lack of learning on the receivers part. Such a lack of learning, in turn, was due to the very little learning shown by the senders at the outset of evolution. So this simulation produced a sort of vicious circle, where little initial learning prevented itself from becoming more substantial across generations.

The main implication of this effect is that in order to obtain a substantial evolution of social behavior, a previous evolutionary process must take place in which the effect in question does not arise. In principle, such a process would have to involve training with a maximally strong CS1, so that any adaptive advantage provided by learning to respond to this CS is favored across generations. This was the main motivation for running the next simulation.

Simulation 2

This simulation was divided into two phases. In Phase 1, a lineage was created from the same founder population used in Simulation 1, with the difference that all networks were trained with CS1. As the upper panel of Figure 1 shows, the mean genetic overlap (measure of genetic similarity) increased substantially from G0 to G50, which indicates that individuals in G50 were genetically far more similar among themselves than individuals in G0. The lower panel shows that mean population fitness also increased substantially from G0 to G50. Also, individuals in the last generation showed an increase in the average *CR/UR* activation across training trials, from close-to-minimum to close-to-maximum values, indicating a substantial learning to respond to CS1. These results reproduce previous ones (see Burgos, 1997). Then, in Phase 2, Simulation 1 was repeated but using G50 as a founder population.

The upper panel of Figure 2 shows that the mean genetic overlap remained roughly the same from G50 to G100 (the last generation of Phase 2). However, this does not necessarily mean that these two populations were genetically similar. Indeed, the lower panel shows two effects. First, the new selection pressure substantially decreased the average fitness for G50, relative to the fitness of the very same generation under the selection pressure in Phase 1. This rather sharp decrease in the average population fitness was due to the fact that the CS2 ISI for the receivers in this phase was much shorter than the CS1 ISI for the senders in Phase 1 and Phase 2. Indeed, while the CS1 ISI had an effective

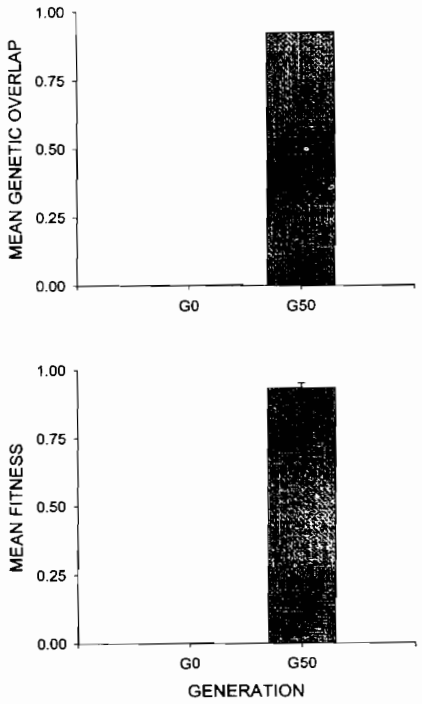

Figure 1. Changes in mean population genetic overlap (upper panel) and mean population fitness (lower panel) from G0 to G50, Phase 1, Simulation 2. Genetic overlap was a measure of genetic similarity between pairs of individuals. An overlap of zero indicates that the two genotypes have the same alleles in half of their genes. This is the expected value for a randomly generated population. An overlap of 1 indicates that the two genotypes have the same alleles in all of their genes. Individual fitness was defined as the mean activation per CR/UR NPE per evaluation trial after training. Error bars represent standard errors.

duration 6 ts, the CS2 ISI lasted approximately between 2 and 4 ts. This happened because CS2 was identical to the CR by the senders to CS1, and this CR tended to have an approximate delay between 1 and 4 ts. That is to say, while CS1 (the nonsocial stimulus given to the senders) involved a constant, maximum input activation across ts, CS2 (the social stimulus given to the receivers) involved a variable input activation.

The second effect was that selection for responding with a high *CR/UR* activation to the behavior of another individual during Phase 2 resulted in a substantial increase of average fitness population. This effect indicates that evolution of this kind of responding occurred, despite the fact that the genetic variation in G50 was

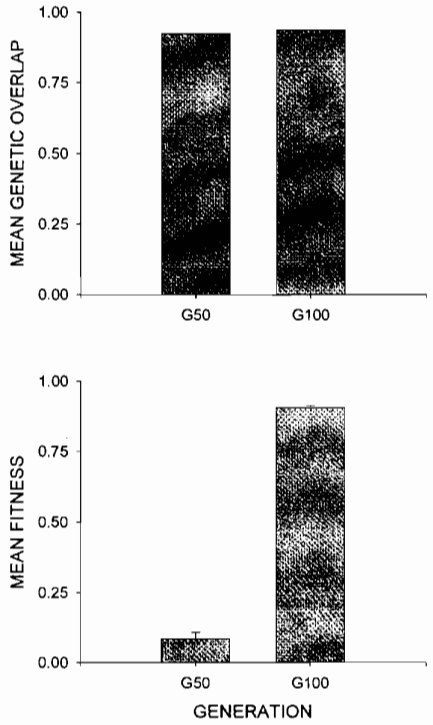

Figure 2. Changes in mean population genetic overlap (upper panel) and mean population fitness (lower panel) from G50 to G100, Phase 2, Simulation 2. Error bars represent standard errors.

very low. If we view G50 and G100 as representing phenotypically (and, hence, genotypically) different species, an implication of this result is that selection for responding to the behavior of another individual might be sufficient to produce speciation, that is, the generation of one species from a different one. As Figure 3 shows, G50 and G100 were phenotypically (and, hence, genotypically) different. In fact, the social contingency caused a decrease in the network size. This effect was due to the fact that the ISI for CS2 was shorter than the ISI for CS1. As mentioned before, previous simulations based on the present approach have shown a systematic relation between network architecture and performance, under temporally different stimulus relations. More specifically, smaller networks tend to perform better under arrangements that involve a shorter ISI than under arrangements that involve a longer ISI (Donahoe & Burgos, in press). Also, selection for high CR under a shorter ISI results in smaller networks than selection

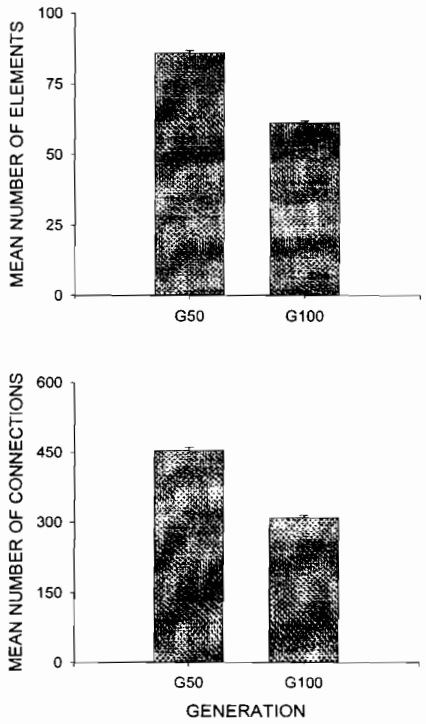

Figure 3. Changes in the mean number of elements (upper panel) and mean number of connections (lower panel) from G50 to G100, Phase 2, Simulation 2. Error bars represent standard errors.

for high CR under longer ISI (Burgos, 1997). The present findings thus replicate these results.

Simulation 3

In the previous simulation, speciation due to selection for responding to the behavior of another individual (Phase 2) was obtained from the last generation of a single lineage (Phase 1). This result, however, may be specific to the particular population obtained in Phase 1 (viz., G50). In order to test the generality of this result, a genetically (and, hence, phenetically) more diverse population should be obtained from the last generation of several different lineages.

For this simulation, then, a total of 10 different lineages were generated, each one with its own random initial population. Each lineage repeated Phase 1 of the previous simulation, leading to comparable results. That is, selection for respond-

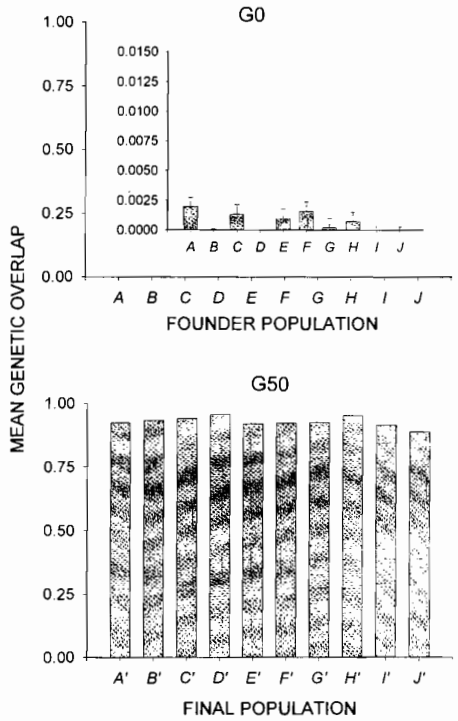

Figure 4. Mean genetic overlap for each population in G0 (upper panel) and G50 (lower panel), Phase 1, Simulation 3. The inset in the upper panel shows mean genetic overlaps for G0 in a different scale. Error bars represent standard errors.

ing with a high *CR/UR* activation to CS1 resulted in a substantial increase in mean population genetic overlap (see Figure 4) and mean population fitness (see Figure 5) from G0 to G50. Then, a sample of ten genotypes was randomly selected from the last generation of each lineage, to obtain a new founder population for Phase 2. As Figure 6 shows, the mean number of elements and number of connections of the networks constituting these samples can be ordered in a scale, analogous to the kind one can construct from certain quantitative properties of the brains of actual species (e.g., see Eccles, 1989, pp. 39-42). Of course, strictly speaking, the present scale is not a phylogenetic one, for the resulting populations came from different, independent, random initial populations. Therefore, they were not phylogenetically related. But still, the model is sufficiently powerful to generate equally fit, although phenetically and genetically different populations, from genetically different initial populations, through identical selection pressures. An implication of this outcome for the phylogeny of behavior in actual biological forms

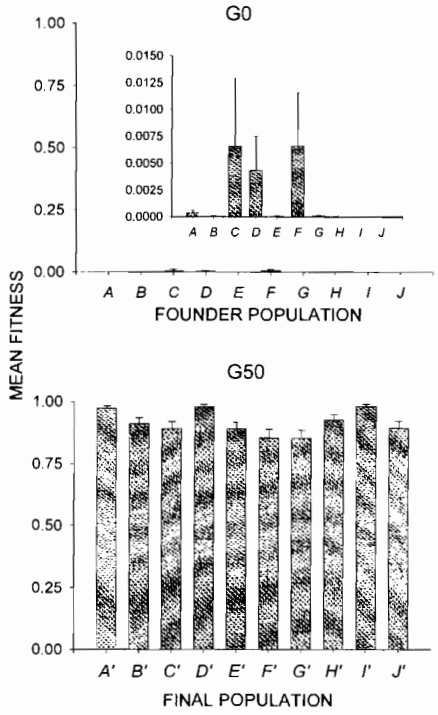

Figure 5. Mean fitness for each population in G0 (upper panel) and G50 (lower panel), Phase 1, Simulation 3. The inset in the upper panel shows mean fitness for G0 in a different scale. Error bars represent standard errors.

is that identical selection pressures, working on genetically different populations, lead to genetically (and, hence, phenetically) different populations.

Then, Phase 2 was run, which was identical to Phase 2 of the previous simulation. As the upper panel of Figure 7 shows, the mean genetic overlap for G50 was below .25, indicating that genetic variation in the population was considerable. But after 50 generations of selection for responding to the behavior of another individual, the mean population genetic overlap increased substantially. The lower panel shows that this selection also resulted in a substantial increase in the mean population fitness from G50 to G100. Note that the average fitness for G50 was considerably higher than the one for G50 in the previous simulation (roughly .35 versus .08, respectively). This outcome is due to the fact that some receivers in the present G50 were able, at the outset, to learn to respond with a relatively high CR to CS2, in spite of the fact that their ancestors had never experienced CS2. This phenomenon can be explained by arguing that the prob-

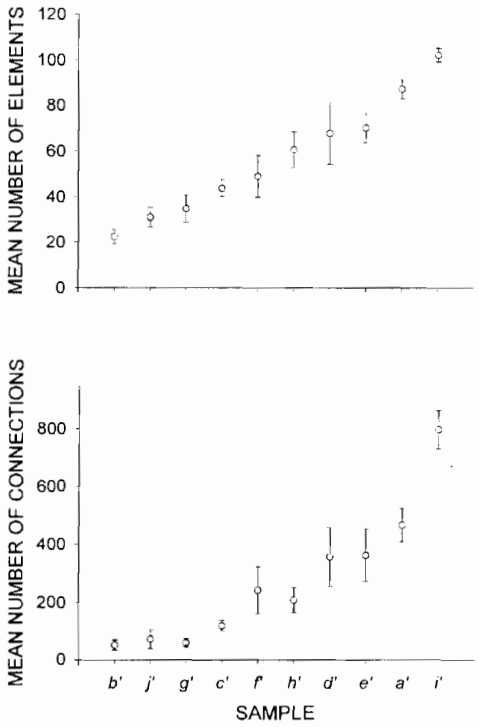

Figure 6. Mean number of elements (upper panel) and connections (lower panel) for random samples (n = 10) of each population in G50. These samples constituted the founder population for Phase 2 of Simulation 3. Samples were ordered according to the mean number of elements. Each sample was identified as the lower-case letter of its population (e.g., a' was a random sample of ten individuals from population A' in G50 of Phase 1). Error bars represent standard deviations.

ability of obtaining such individuals after generating ten lineages (the case in the present simulation) is substantially larger than after generating only one lineage (the case of Simulation 2).

Finally, as Figure 8 shows, selection for responding to the behavior of another individual resulted in genetically and phenetically very similar individuals. Such a drastic reduction in variation is due to the fact that selection acted upon a single initial population. Also, the last generation consisted on rather small networks. This outcome can be explained, once again, by the fact that the environmental arrangement for the receivers involved a shorter ISI than the one used in training during Phase 1 (see Figure 3, about Phase 2 of Simulation 2).

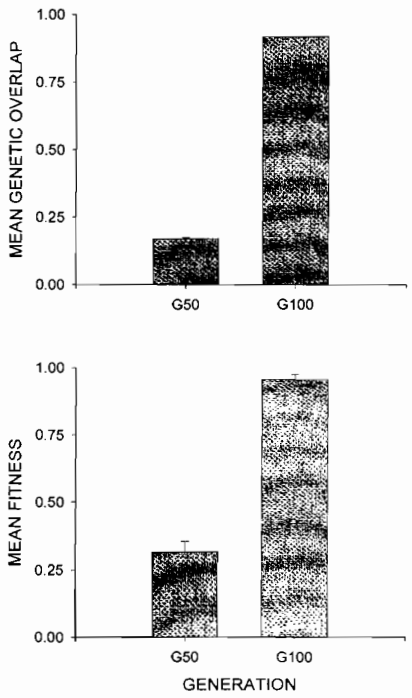

Figure 7. Changes in mean population genetic overlap (upper level) and mean population fitness from G50 to G100, Phase 2, Simulation 3. G50 consisted of random samples from the different populations obtained in Phase 1. Error bars represent standard errors.

DISCUSSION

I have presented an approach to behavior in general and to social behavior in particular, based on the key notion that an understanding of social behavior requires a synthesis of biological and behavioral models that have been obtained through conceptual and experimental analysis. The proposed synthesis assumes a conceptualization of biobehavioral phenomena as being organized into a hierarchy of levels of analysis or organization, which roughly correspond to those identified in psychology and biology. Thus we have that the level of the NPE corresponds to the cellular level, the level of the network to the neuroanatomical level, and the of the network behaving in time under certain environmental conditions to the level of behavior. An additional, higher level arises from populations of individuals across time, related through descent, which corresponds to the level of phylogeny and evolution by selection. The approach is nonreductive in

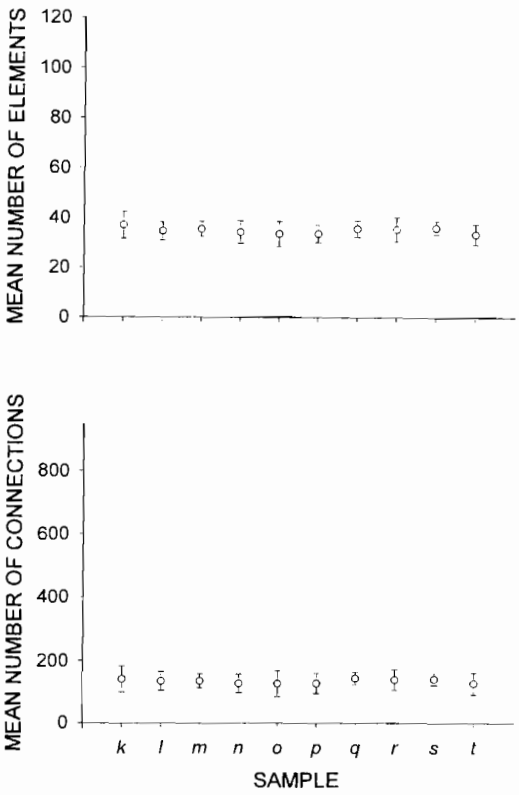

Figure 8. Mean number of elements (upper panel) and connections (lower panel) for random samples (n = 10) of each population in G100, Phase 2, Simulation 3. Error bars represent standard deviations.

that it does not allow for a complete deduction of processes at one level from processes at a lower level.

The synthesis was given a computational form, which allowed for empirical studies through computer simulations. The simulation results suggest that social behavior is not fundamentally different from (i.e., follows the same underlying biobehavioral principles as) any other kind of behavior, at least with respect to the aspect studied (viz., responding by one individual to the behavior of another). However, social behavior may be unique regarding the particular realization of the underlying principles. Indeed, the present results suggest that social behavior requires a previous evolutionary process that is based on selection for nonsocial behavior. Learning to respond to nonsocial stimuli may very well evolutionarily

precede learning to respond to social stimuli. Also, ontogenetically speaking, social contingencies may be unique, at least with respect to temporal relations among stimuli. Specifically, if an individual responds in a delayed fashion to a certain, nonsocial stimulus, and such a response serves as a stimulus to another individual, then both individuals are under the control of different temporal relations, which, in turn, may make them behave differently.

In the present simulations, that difference was shown by the fact that the response of the senders to CS1 in Phase 2 of Simulation 3 was delayed, for which it was necessarily shorter than CS1. When this response served as a CS2 for the receivers, it was equally shorter. The nonsocial stimulus for the senders (CS1) thus was necessarily shorter than the social stimulus for the receivers (CS2), which made the behavior of the two individuals to become under the control of different temporal relations.

Regarding the particular kind of dyad studied here (the one-directional sender-receiver dyad) the simulations showed that an individual required a previous nonsocial experience before it can effectively function as a sender. An implication of this finding for analogous dyads found in actual biological forms (e.g., predator-prey, leader-follower, or teacher-student dyads) is that the leading component of the dyad must at least have had some previous experience with nonsocial circumstances. Hence, what applies to phylogeny also applies to ontogeny, in that nonsocial learning is the foundation of social learning.

Most social situations in actual biological forms, of course, are far more complex than the one given by the basic sender-receiver dyad. In fact, many instances of this kind of dyad may not be social at all, under certain definitions of social behavior. In any case, a slightly more complex variant of this kind of dyad would be one in which the roles of sender and receiver are more dynamic, more changing within a generation. For example, a pair of individuals may interact in such a way that they play different functions at different moments in time. The specific form that this variant can take in the present approach involves training first both individuals to respond to CS1 (the nonsocial stimulus). Then, they would play the different roles (either sender or receiver) in different trials, so that eventually both learn to respond to CS2 (the social stimulus). A preliminary simulation in this direction was run, but the results showed that once networks learned to respond to CS1, they were unable to learn to respond to CS2. That is, networks showed a severe form of synaptic blocking, one that did not even require the concurrent presence of the previously trained CS. This kind of result is due to the fact that the learning rule of the neurocomputational model includes competition among presynaptic inputs for a limited amount of connection weight. This characteristic of the model was very likely amplified by the fact that individuals were selected for responding to only one CS, which resulted in populations of specialists (i.e., individuals that could learn only one thing at a time). So, perhaps, explicit selection for responding to two different nonsocial CSs may be required before a more interactive kind of behavioral relation can be simulated.

REFERENCES

Adami, C., Belew, R. K., Kitano, H., & Taylor, C. E. (1998). *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*. Cambridge, MA: Bradford.

Bickle, J. (1998). *Psychoneural reduction: The new wave*. Cambridge, MA: Bradford.

Brooks, R. A., & Maes, P. (1994). *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of living systems*. Cambridge, MA: The MIT Press.

Brown, J. R. (1991). *The laboratory of the mind: Thought experiments in the natural sciences*. London: Routledge.

Brown, M. C., Hopkins, W. G., & Keynes, R. J. (1991). *Essentials of neural development*. Cambridge University Press.

Burgos, J. E. (1995). *Computational explorations of the evolution of the evolution of artificial neural networks in Pavlovian environments*. Unpublished Doctoral Dissertation, Neuroscience and Behavior Program, University of Massachusetts/Amherst.

Burgos, J. E. (1997). Evolving artificial neural networks in Pavlovian environments. In J. W.

Donahoe and V. Packard-Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations* (pp. 58-79). New York: Elsevier.

Burgos, J. E. (1999). Una reconstrución neurocomputacional del problema de los dos tipos de condicionamiento. In A. L. Rangel, L. M. Sánchez, M. Lozada, & C. Silva (Eds.), *Contribuciones a la psicología en Venezuela* (pp. 215-241). Universidad Central de Venezuela.

Burgos, J. E., & Donahoe, J. W. (2000). Structure and function in selectionism: Implications for complex behavior. In J. C. Leslie and D. Blackman (Eds.), *Experimental and applied analyses of human behavior* (pp. 39-57). Reno, NV: Context Press.

Darwin, C. (1859). *The origin of species by means of natural selection*. New York: Avenel.

Dawkins, R. (1982). *The extended phenotype*. Oxford University Press.

Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon & Schuster.

Donahoe, J. W., Burgos, J. E., & Palmer, D. C. (1993). A selectionist approacy to reinforcement. *Journal of the Experimental Analysis of Behavior*, **60**, 17-40.

Donahoe, J. W., Crowley, M. A., Millard, W. J., & Stickney, K. A. (1982). A unified principle of reinforcement: Some implications for matching. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts* (pp. 493-521). Cambridge, MA: Ballinger.

Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior*. Boston: Allyn & Bacon.

Donahoe, J. W., Palmer, D. C., & Burgos, J. E. (1997a). The S-R issue: Its status in behavior analysis and in Donahoe and Palmer's *Learning and complex behavior*. *Journal of the Experimental Analysis of Behavior*, **67**, 193-211.

Donahoe, J. W., Palmer, D. C., & Burgos, J. E. (1997b). The unit of selection: What do reinforcers reinforce? *Journal of the Experimental Analysis of Behavior*, **67**, 259-273.

Eccles, J. C. (1989). *Evolution of the brain: Creation of the self*. London: Routledge.

Emmeche, C. (1991). *The garden in the machine: The emerging science of Artificial Life*. Translated by S. Sampson (1994). Princeton, New Jersey: Princeton University Press.

Fukushima, K. (1975). Cognitron: A self-organizing multilayered neural network. *Biological Cybernetics*, **20**, 121-136.

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.

Gould, S. J. (1989). *Wonderful life: The Burgess Shale and the nature of history*. New York: Norton.

Holland, J. H. (1992). *Adaptation in natural and artificial systems*. Cambridge, MA: MIT Press.

Hollis, K. L. (1984). The biological function of of Pavlovian conditioning: The best defense is a good offense. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 413-425.

Hollis, K. L. (1990). The role of Pavlovian conditioning in territorial aggression and reproduction. In D. A. Dewsbury (Ed.), *Contemporary issues in comparative psychology* (pp. 197-219). Sunderland, MA: Sinauer.

Kamil, A. C. (1994). A synthetic approach to the study of animal intelligence. In L. A. Real (Ed.), *Behavioral mechanisms in evolutionary ecology* (pp. 11-45). The University of Chicago Press.

Kandel, E. R., & Schwartz, J. H. (1982). Molecular biology of learning: Modulation of transmitter release. *Science*, **218**, 433-443.

Langton, C. G. (1992). *Artificial Life II*. Redwood City, CA: Addison-Wesley.

Langton, C. G. (1994). *Artificial Life III*. Redwood City, CA: Addison-Wesley. Langton, C. G., & Shimohara, T. (1997). *Artificial Life V: Proceedings of th Fifth International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, MA: MIT Press.

Mitchell, M. (1996). *An introduction to genetic algorithms*. Canbridge, MA: MIT Press.

Teitelbaum, P. (1977). Levels of integration of the operant. In W. K. Honig and J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 7-27). Englewood Cliffs, NJ: Prentice-Hall.

Wilson, E. (1975). *Sociobiology: The new synthesis.* Harvard University Press.