

In Mark Johnson (Ed.), *Brain development and cognition: A reader* (2nd ed.).
Oxford: Blackwell Publishers, 2002.

CONNECTIONISM AND THE STUDY OF CHANGE

Elizabeth A. Bates

Jeffrey L. Elman

University of California, San Diego

Developmental psychology and developmental neuropsychology have traditionally focused on the study of children. But these two fields are also supposed to be about the study of change, i.e. changes in behavior, changes in the neural structures that underlie behavior, and changes in the relationship between mind and brain across the course of development. Ironically, there has been relatively little interest in the mechanisms responsible for change in the last 15–20 years of developmental research. The reasons for this de-emphasis on change have a great deal to do with a metaphor for mind and brain that has influenced most of experimental psychology, cognitive science and neuropsychology for the last few decades, i.e. the metaphor of the serial digital computer. We will refer to this particular framework for the study of mind as the *First Computer Metaphor*, to be contrasted with a new computer metaphor, variously known as connectionism, parallel distributed processing, and/or neural networks. In this brief chapter, we will argue that the First Computer Metaphor has had some particularly unhappy consequences for the study of mental and neural development. By contrast, the *Second Computer Metaphor* (despite its current and no doubt future limitations) offers some compelling advantages for the study of change, at both the mental and the neural level.

The chapter is organized as follows: (1) a brief discussion of the way that change has (or has not) been treated in the last decade of research in developmental psychology, (2) a discussion of the First Computer Metaphor, and its implications for developmental research, (3) an introduction to the Second Computer Metaphor, and the promise it offers for research on the development of mind and brain, ending with (4) a response to some common misconceptions about connectionism.

1. What Happened to the Study of Change?

Traditionally, there are three terms that have been used to describe changes in child behavior over time: maturation, learning, and development. For our purposes here, these terms can be defined as follows.

(a) *Maturation*. As the term is typically used in the psychological literature (although this use may not be entirely accurate from a biological perspective—Bates, Thal, Finlay and Clancy, in press; Elman et al., 1996; Johnson, 1997), “maturation” refers to the timely appearance or unfolding of behaviors that are predetermined, in their structure and their sequence, by a well-

defined genetic program. The role of experience in a strong maturational theory is limited to a “triggering” function (providing the general or specific conditions that allow some predetermined structures to emerge) or a “blocking” function (providing conditions that inhibit the expression of some predetermined event). The environment does not, in and of itself, provide or cause behavioral structure.

(b) *Learning*. “Learning” is typically defined as a systematic change in behavior as a result of experience. Under some interpretations, learning refers to a copying or transfer of structure from the environment to the organism (as in “acquisition” or “internalization”). Under a somewhat weaker interpretation, learning may refer to a shaping or alteration of behavior that is caused by experience, although the resulting behavior does not resemble structures in the environment in any direct or interesting way.

(c) *Development*. As defined by Werner (1948) in his elaboration of the “orthogenetic principle”, “development” refers to any positive change in the internal structure of a system, where “positive” is further defined as an increase in the number of internal parts (i.e. differentiation), accompanied by an increase in the amount of organization that holds among those parts. Under this definition, the term “development” is neutral to the genetic or experiential sources of change, and may include emergent forms that are not directly predictable from genes or experience considered separately (i.e. the sum is greater than and qualitatively different from the parts).

Although all three terms have been used to describe behavioral change in the psychological literature, the most difficult and (in our view) most interesting proposals are the ones that have involved emergent form, i.e. changes that are only indirectly related to structure in the genes or the environment (Bates et al., 1998; Elman et al., 1996; MacWhinney, 1999). We are referring here not to the banal interactionism in which black and white yield grey, but to a much more challenging interactionism in which black and white converge and interact to yield an unexpected red. Because this interactionist view appears to be the only way to explain how new structures arise, it may be our only way out of a fruitless nature/nurture debate that has hampered progress in developmental psychology for most of its history.

Within our field, the most complete interactionist theory of behavioral change to date is the theory offered by Jean Piaget, across a career that spanned more than

fifty years (Piaget, 1952, 1970a, 1970b, 1971). Piaget's genetic epistemology concentrated on the way that new mental structures emerge at the interface between an active child and a structured world. The key mechanism for change in Piaget's theory is the consummate biological notion of adaptation. Starting with a restricted set of sensorimotor schemes (i.e. structured "packages" of perception and action that permit activities like sucking, reaching, tracking, and/or grasping), the child begins to act upon the world (assimilation). Actions are modified in response to feedback from that world (accommodation), and in response to the degree of internal coherence or stability that action schemes bear to one another (reciprocal assimilation). The proximal cause that brings about adaptation is a rather poorly defined notion of equilibration, i.e. the re-establishment of a stable and coherent state after a perturbation that created instability or disequilibrium. In the infant years, adaptation of simple sensorimotor schemes to a structured world leads to an increasingly complex and integrated set of schemes or "plans", structures that eventually permit the child to "re-present" the world (i.e. to call potential perceptuo-motor schemes associated with a given object or event into an organized state-of-readiness, in the absence of direct perceptual input from the represented object or event). This developmental notion of representation comprised Piaget's explanation for the appearance of mental imagery, language and other symbolic or representational forms somewhere in the second year of life. After this point, the process of adaptation continues at both the physical and representation level (i.e. operations on the real world, and operations on the new "mental world"), passing through a series of semi-stable "stages" or moments of system-wide equilibrium, ultimately leading to our human capacity for higher forms of logical reasoning.

This "bootstrapping" approach to cognitive development does involve a weak form of learning (as defined above), but the mental structures that characterize each stage of development are not predictable in any direct way from either the structure of the world or the set of innate sensorimotor schemes with which the child began. Furthermore, Piaget insisted that these progressive increases in complexity were a result of activity ("construction"), and not a gradual unfolding of predetermined forms (maturation). In this fashion, Piaget strove to save us from the nature/nurture dilemma. Behavioral outcomes were determined not only by genes, or by environment, but by the mathematical, physical and biological laws that determine the kinds of solutions that are possible for any given problem. As Piaget once stated in a criticism of his American colleague Noam Chomsky, "That which is inevitable does not have to be innate" (Piaget, 1970a).

There was a period in the history of developmental psychology in which Piagetian theory assumed a degree of orthodoxy that many found stifling. Decades later, it now appears that much of Piaget's theory was wrong in detail. For one thing, it is now clear that the infant's initial stock of innate sensorimotor schemes is far richer

than Piaget believed. It is also clear that Piaget overestimated the degree of cross-domain stability that children are likely to display at any given point in development (i.e. the notion of a coherent "stage"). Once the details of his stage theory were proven inadequate, all that really remained were the principles of change that formed the bedrock of Piaget's genetic epistemology— notions of adaptation and equilibration that struck many of his critics as hopelessly vague, and a notion of emergent form that many found downright mystical. Piaget was aware of these problems, and spent the latter part of his career seeking a set of formalisms to concretize his deep insights about change. Most critics agree that these efforts failed. This failure, coupled with new empirical information showing that many other aspects of the theory were incorrect, has led to a widespread repudiation of Piaget. Indeed, we are in a period of "anti-Piagetianism" of patricidal dimensions.

But what have we put in Piaget's place? We have never replaced his theory with a better account of the epistemology of change. In fact, the most influential developmental movements of the last two decades have essentially disavowed change. Alas, we fear that we are back on the horns of the nature-nurture dilemma from which Piaget tried in vain to save us.

On the one hand, we have seen a series of strong nativist proposals in the last few years, including proposals by some neo-Gibsonian theorists within the so-called "competent infant movement" (Baillargeon, 1999; Baillargeon and de Vos, 1991; Spelke, 1990, 1991; Spelke and Newport, 1998), and proposals within language acquisition inspired by Chomsky's approach to the nature and origins of grammar (Hyams, 1986; Lightfoot, 1991; Roeper and Williams, 1987). In both these movements, it is assumed that the essence of what it means to be human is genetically predetermined. Change—insofar as we see change at all—is attributed to the maturation of predetermined mental content, to the release of preformed material by an environmental "trigger", and/or to the gradual removal of banal sensory and motor limitations that hid all this complex innate knowledge from view. Indeed, the term "learning" has taken on such negative connotations in some quarters that efforts are underway to eliminate it altogether. The following quotes from Piatelli-Palmarini (1989) illustrate how far things have gone:

"I, for one, see no advantage in the preservation of the term learning. We agree with those who maintain that we would gain in clarity if the scientific use of the term were simply discontinued." (p. 2)

"Problem-solving... adaptation, simplicity, compensation, equilibration, minimal disturbance and all those universal, parsimony-driven forces of which the natural sciences are so fond, recede into the background. They are either scaled down, at the physico-chemical level, where they still make a lot of sense, or dismissed altogether." (pp. 13-14)

On the other hand, the neo-Vygotskian movement and associated approaches to the social bases of cognition have provided us with another form of pre-

formationalism, insisting that the essence of what it means to be human is laid out for the child in the structure of social interactions (Bruner and Sherwood, 1976; Rogoff, 1990). In these theories, change is viewed primarily as a process of internalization, as the child takes in preformed solutions to problems that lie in the “zone of proximal development”, i.e. in joint activities that are just outside his current ability to act alone. Related ideas are often found in research on “motherese”, i.e. on the special, simplified and caricatured form of language that adults direct to small children (for a review, see Ferguson and Snow, 1978). In citing these examples, we do not want to deny that society has an influence on development, because we are quite sure that it does. Our point is, simply, that the pendulum has swung too far from the study of child-initiated change. The most influential movements in developmental psychology for the last two decades are those that have de-emphasized change in favor of an emphasis on some kind of preformation: either a preformation by nature and the hand of God, or a preformation by the competent adult.

Why have we accepted these limits? Why haven't we moved on to study the process by which new structures really do emerge? We believe that developmental psychology has been influenced for many years by a metaphor for mind in which it is difficult to think about change in any interesting form—which brings us to the First Computer Metaphor.

2. The First Computer Metaphor and its Implications for Development

At its core, the serial digital computer is a machine that manipulates symbols. It takes individual symbols (or strings of symbols) as its input, applies a set of stored algorithms (a program) to that input, and produces more symbols (or strings of symbols) as its output. These steps are performed one at a time (albeit very quickly) by a central processor. Because of this serial constraint, problems to be solved by the First Computer must be broken down into a hierarchical structure that permits the machine to reach solutions with maximum efficiency (e.g. moving down a decision tree until a particular subproblem is solved, and then back up again to the next step in the program).

Without question, exploitation of this machine has led to huge advances in virtually every area of science, industry and education. After all, computers can do things that human beings simply cannot do, permitting quantitative advances in information processing and numerical analysis that were unthinkable a century ago. The problem with this device for our purposes here lies not in its utility as a scientific tool, but in its utility as a scientific metaphor, in particular as a metaphor for the human mind/brain. Four properties of the serial digital computer have had particularly unfortunate consequences for the way that we have come to think about mental and neural development.

(1) *Discrete representations.* The symbols that are manipulated by a serial digital computer are discrete entities. That is, they either are or are not present in the input. There is no such thing as 50% of the letter **A** or 99% of the number **7**. For example, if a would-be user types in a password that is off by only one keystroke, the computer does not respond with “What the heck, that's close enough.” Instead, the user is damned just as thoroughly as he would be if he did not know the password at all.

People (particularly children) rarely behave like this. We can respond to partial information (degraded input) in a systematic way; and we often transform our inputs (systematic or not) into partial decisions and imperfect acts (degraded output). We are error-prone, but we are also forgiving, flexible, willing and able to make the best of what we have. This mismatch between human behavior and the representations manipulated by serial digital computers has of course been known for some time. To resolve this well-known discrepancy, the usual device adopted by proponents of the First Computer Metaphor for Mind is the competence/performance distinction. That is, it is argued that our knowledge (competence) takes a discrete and idealized form that is compatible with the computer metaphor, but our behavior (performance) is degraded by processing factors and other sources of noise that are irrelevant to a characterization of knowledge and (by extension) acquisition of knowledge. This is a perfectly reasonable intellectual move, but as we will see in more detail below, it has led to certain difficulties in characterizing the nature of learning that often result in the statement that learning is impossible.

(2) *Absolute rules.* Like the symbolic representations described above, the algorithms contained in a computer program also take a discrete form. If the discrete symbols that trigger a given rule are present in the input, then that rule must apply, and give an equally discrete symbol or string of symbols as its output. Conversely, if the relevant symbols are not present in the input, then the rule in question will not apply. There is no room for anything in between, no coherent way of talking about 50% of a rule, or (for that matter) weak vs. strong rules. Indeed, this is exactly the reason why computers are so much more reliable than human beings for many computational purposes.

Presented with the well-known mismatch between human behavior and the absolute status of rules in a serial digital computer, proponents of the First Computer Metaphor for Mind usually resort to the same competence/performance described above. Alternatively, there have been attempts to model the probabilistic nature of human behavior by adding weights to rules, a device that permits the model to decide which rule to apply (or in what order of preference) when a choice has to be made. The problem is that these weights are in no way a natural product or property of the architecture in which they are embedded, nor are they produced automatically by the learning process. Instead, these weights

are arbitrary, *ad hoc* devices that must be placed in the system by hand—which brings us to the next point.

(3) *Learning as programming.* The serial digital computer is not a self-organizing system. It does not learn easily. Indeed, the easiest metaphor for learning in a system of this kind is programming; that is, the rules that must be applied to inputs of some kind are placed directly into the system—by man, by Nature or by the hand of God. To be sure, there is a literature on computer learning in the field of artificial intelligence. However, most of these efforts are based on a process of hypothesis testing. In such learning models, two essential factors are provided *a priori*: a set of hypotheses that will be tested against the data, and an algorithm for deciding which hypothesis provides the best fit to those data. This is by its very nature a strong nativist approach to learning. It is not surprising that learning theories of this kind are regularly invoked by linguists and psycholinguists with a strong nativist orientation. There is no graceful way for the system to derive new hypotheses (as opposed to modifications of a pre-existing option). Everything that really counts is already there at the beginning.

Once again, however, we have an unfortunate mismatch between theory and data in cognitive science. Because the hypotheses tested by a traditional computer-learning model are discrete in nature (based on the rule and representations described above), learning (a.k.a. “selection”) necessarily involves a series of discrete decisions about the truth or falsity of each hypothesis. Hence we would expect change to take place in a crisp, step-wise fashion, as decisions are made, hypotheses are discarded, and new ones are put in their place. But human learning rarely proceeds in this fashion, characterized more often by error, vacillation and backsliding. In fact, the limited value of the serial digital computer as a metaphor for learning is well known. Perhaps for this reason, learning and development have receded into the background in modern cognitive psychology, while the field has concentrated instead on issues like the nature of representation, processes of recognition and retrieval, and the various stages through which discrete bits of information are processed (e.g. various buffers and checkpoints in a serial process of symbol manipulation). Developmental psychologists working within this framework (or indirectly influenced by it) have moved away from the study of change and self-organization toward a catalogue of those representations that are there at the beginning (e.g. the “competent infant” movement in cognition and perception; the parameter-setting movement in developmental psycholinguistics), and/or a characterization of how the processes that elaborate information mature or expand across the childhood years (i.e. changes in performance that “release” the expression of pre-existing knowledge).

(4) *The hardware/software distinction.* One of the most unfortunate consequences of the First Computer Metaphor for cognitive science in general and developmental psychology in particular has been the acceptance

of a strong separation between software (the knowledge—symbols, rules, hypotheses, etc.—that is contained in a program) and hardware (the machine that is used to implement that program). From this perspective, the machine itself places very few constraints on our theory of knowledge and (by extension) behavior, except perhaps for some relatively banal concerns about capacity (e.g. there are some programs that one simply cannot run on a small personal computer with limited memory).

The distinction between hardware and software has provided much of the ammunition for an approach to philosophy of mind and cognitive science called Functionalism (Fodor, 1981).¹ Within the functionalist school, the essential properties of mind are derived entirely from the domains on which the mind must operate: language, logic, mathematics, three-dimensional space, etc. To be sure, these properties have to be implemented in a machine of some kind, but the machine itself does not place interesting constraints on mental representations (i.e. the objects manipulated by the mind) or functional architecture (i.e. the abstract system that manipulates those objects). This belief has justified an approach to cognition that is entirely independent of neuroscience, thereby reducing the number and range of constraints to which our cognitive theories must respond. As a by-product (since divorces usually affect both parties), this approach has also reduced the impact of cognitive theories and cognitive phenomena on the field of neuroscience.

The separation between biology and cognition has had particularly serious consequences for developmental psychology, a field in which biology has traditionally played a major role (i.e. a tradition that includes Freud, Gesell, Baldwin, and Piaget, to name a few). Not only have we turned away from our traditional emphasis on change, but we have also turned away from the healthy and regular use of biological constraints on the study of developing minds. Ironically, some of the strongest claims about innateness in the current literature have been put forth in complete disregard of biological facts (Bates et al., in press; Elman et al., 1996). For example, claims about the innateness of detailed cortical representations (e.g. an innate theory of physics, or innate theory of mind) are difficult to square with what we now know about (a) the plastic and activity-dependent nature of the processes that set up cortical microcircuitry, and (b) the mathematical bottleneck involved in setting up a brain with 10^{14} connections using fewer than 10^6 genes, 98% of which humans share with chimpanzees (Gerhart & Kirschner, 1998). The underlying assumption appears to be that our cognitive findings have priority, and if there is a mismatch between

¹ This particular school of Functionalism has little to do with, and is indeed diametrically opposed to, an approach within linguistics and psycholinguistics alternatively called Functional Grammar or Cognitive Linguistics. For discussions, see Bates and MacWhinney, 1989; Givón, 1984; Lakoff, 1987; Langacker, 1987.

cognitive and biological conclusions, we probably got the biology wrong (which may be the case some of the time—but surely not all the time!).

It seems to us that we need all the constraints that can be found to make sense of a growing mass of information about cognitive development, language development, perceptual development, social development. Furthermore, we suspect that developmental neuroscience would also profit from a healthy dose of knowledge about the behavioral functions of the neural systems under study. Finally, we would all be better off if we could find a computational model (or class of models) in which it would be easier to organize and study the mutual constraints that hold between mental and neural development—which brings us to the next computer metaphor.

3. The Second Computer Metaphor and Its Implications for Development

During the 1950s and 60s, when the First Computer Metaphor for mind began to influence psychological research, some information scientists were exploring the properties of a different and competing computational device called the Perceptron (Rosenblatt, 1958, 1962). The roots of this approach can be traced to earlier work in cybernetics (Minsky, 1956; von Neumann, 1951, 1958) and in neurophysiology (Eccles, 1953; Hebb, 1949; McCulloch and Pitts, 1943). In a perceptron network, unlike the serial digital computer, there was not a clear distinction between processor and memory, nor did it operate on symbols in the usual sense of the term. Instead, the perceptron network was composed of a large number of relatively simple “local” units that worked in parallel to perceive, recognize and/or categorize an input. These local units or “nodes” were organized into two layers, an “input set” and an “output set”. In the typical perceptron architecture, every unit on the input layer was connected by a single link to each and every unit on the output layer (see Figure 1). These connections varied in degree or strength, from 0 to 1 (in a purely excitatory system) or from -1 to $+1$ (in a system with both activation and inhibition). A given output unit would “fire” as a function of the amount of input that it received from the various input units, with activation collected until a critical firing threshold was reached (see also McCulloch and Pitts, 1943). Individual acts of recognition or categorization in a Perceptron reflect the collective activity of all these units. Knowledge is a property of the connection strengths that hold between the respective input and output layers; the machine can be said to “know” a pattern when it gives the correct output for a given class of inputs (including novel members of the input class that it has never seen before, i.e. generalization).

There are some obvious analogies between this system and the form of computation carried out in real neural systems, e.g. excitatory and inhibitory links, summation of activation, firing thresholds, and above all the distribution of patterns across a large number of interconnected units. But this was not the only ad-

vantage that perceptrons offered, compared with their competitors. The most important property of perceptrons was (and is) their ability to learn by example.

During the teaching and learning phase, a stimulus is registered on the input layer in a distributed fashion, by turning units on or off to varying degrees. The system produces the output that it currently prefers (based, in the most extreme *tabula rasa* case, on a random set of connections). Each unit in this distributed but “ignorant” output is then compared with the corresponding unit in the “correct” output. If a given output unit within a distributed pattern has “the right answer”, its connection strengths are left unchanged. If a given output has “the wrong answer”, the size of the error is calculated by a simple difference score (i.e. “delta”). All of the connections to that erroneous output are then increased or decreased in proportion to the amount of error that they were responsible for on that trial. This procedure then continues in a similar fashion for other trials. Because the network is required to find a single set of connection weights which allow it to respond correctly to all of the patterns it has seen, it typically succeeds only by discovering the underlying generalizations which relate inputs to outputs. The important and interesting result is that the network is then able to respond appropriately not only to stimuli it has seen before, but to novel stimuli as well. The learning procedure is thus an example of learning inductively.

Compared with the cumbersome hypothesis-testing procedures that constitute learning in serial digital computers, learning really appears to be a natural property of the perceptron. Indeed, perceptrons are able to master a broad range of patterns, with realistic generalization to new inputs as a function of their similarity to the initial learning set. The initial success of these artificial systems had some impact on theories of pattern recognition in humans. The most noteworthy example is Selfridge’s “Pandemonium Model” (Selfridge, 1958), in which simple local feature detectors or “demons” work in parallel to recognize a complex pattern. Each demon scans the input for evidence of its preferred feature; depending on its degree of certainty that the relevant feature has appeared, each demon “shouts” or “whispers” its results. In the Pandemonium Model (as in the Perceptron), there is no final arbiter, no *homunculus* or central executive who puts all these daemonical inputs together. Rather, the “solution” is an emergent property of the system as a whole, a global pattern produced by independent, local computations. This also means that results or solutions can vary in their degree of resemblance to the “right” answer, capturing the rather fuzzy properties of human categorization that are so elusive in psychological models inspired by the serial digital computer.

So far so good. And yet this promising line of research came to a virtual end in 1969, when Minsky and Papert published their famous book *Perceptrons*. Minsky and Papert (who were initial enthusiasts and pioneers in perceptron research) were able to prove that perceptrons are only capable of learning a limited class

of first-order, linearly separable patterns. These systems are incapable of learning second-order relations like “A or B but not both” (i.e. logical exclusive OR), and by extension, any pattern of equivalent or greater complexity and interdependence. This fatal flaw is a direct product of the fact that perceptrons are two-layered systems, with a single direct link between each input and output unit. If A and B are both “on” in the input layer, then they each automatically “turn on” their collaborators on the output layer. There is simply no place in the system to record the fact that A and B are both on simultaneously, and hence no way to “warn” their various collaborators that they should shut up on this particular trial. It was clear even in 1969 that this problem could be addressed by adding another layer somewhere in the middle, a set of units capable of recording the fact that A and B are both on simultaneously, and therefore capable of inhibiting output nodes that would normally turn on in the presence of either A or B. So why not add a set of “in between” units, creating three- or four- or N -layered perceptrons? Unfortunately, the learning rules available at that time (e.g. the simple delta rule) did not work with multilayered systems. Furthermore, Minsky and Papert offered the conjecture that such a learning rule would prove impossible in principle, due to the combinatorial complexity of delta calculations and “distribution of blame” in an N -layered system. As it turns out, this conjecture was wrong (after all, a conjecture is not a proof). Nevertheless, it was very influential. Interest in the perceptron as a model of complex mental processes dwindled in many quarters. From 1970 on, most of artificial intelligence research abandoned this architecture in favor of the fast, flexible and highly programmable serial digital computer. And most of cognitive psychology followed suit. (For a somewhat different account of this history, see Papert, 1988; a good collection of historically important documents can be found in Anderson and Rosenfeld, 1989. A good tutorial on how to do connectionist simulations is Plunkett and Elman, 1997, and Golden, 1996, provides an excellent technical overview of neural networks.)

Parallel distributed processing was revived in the late 1970s and early 1980s, for a variety of reasons. In fact, the computational advantages of such systems were never entirely forgotten (Anderson, 1972; Feldman and Ballard, 1980; Hinton and Anderson, 1981; Kohonen, 1977; Willshaw et al., 1969), and their resemblance to real neural systems continued to exert some appeal (Grossberg, 1968, 1972, 1987). But the current “boom” in parallel distributed processing or “connectionism” was inspired in large measure by the discovery of a learning rule that worked for multilayered systems (Le Cun, 1985; Rumelhart et al., 1986a). The Minsky-Papert conjecture was overturned, and there are now many impressive demonstrations of learning in multilayered neural nets, including learning of N -order dependencies like “A or B but not both” (Rumelhart et al.,

1986b).² Multilayer networks have been shown to be universal function approximators, which means they can approximate any function to any arbitrary degree of precision (Hornik, Stinchcombe and White, 1989). Such a network is shown in Figure 2.

Another reason for the current popularity of connectionism derives from technical advances in the design of parallel computing systems. It has become increasingly clear to computer scientists that we are close to the absolute physical limits on speed and efficiency in serial systems—and yet the largest and fastest serial computers still cannot come close to the speed with which our small, slow, energy-efficient brains recognize patterns and decide where and how to move. As Carver Mead has pointed out (Mead, 1989), it is time to “reverse-engineer Nature”, to figure out the principles by which real brains compute information. It is still the case that most connectionist simulations are actually carried out on serial digital computers (which mimic parallelism by carrying out a set of would-be parallel computations in a series, and waiting for the result until the next wave of would-be parallel computations is ready to go). But new, truly parallel architectures are coming on line (e.g. the now-famous Connection Machine) to implement those discoveries that have been made with pseudo-parallel simulations. Parallel distributed processing appears to be the solution elected by Evolution, and (if Mead is right) computer science will have to move in this direction to capture the kinds of processing that human beings do so well.

For developmental psychologists, the Second Computer Metaphor holds some clear advantages for the study of change in human beings. The first set involves the same four areas in which the First Computer Metaphor has let us down: the nature of representations, rules or “mappings”, learning, and the hardware/software issue. The last two are advantages peculiar to connectionist networks: nonlinear dynamics, and emergent form.

(1) *Distributed representations.* The representations employed in connectionist nets differ radically from the symbols manipulated by serial digital computers. First, these representations are “coarse-coded”, distributed across many different units. Because of this property, it is reasonable to talk about the degree to which a representation is active or the amount of a representation that is currently available in this system (i.e. 50% of an “A” or 99% of the number “7”). This also means that patterns can be built up or torn down in bits and pieces, accounting for the graded nature of learning in most instances, and for the gradual or graded patterns of breakdown that are typically displayed by brain-damaged individuals (Hinton and Shallice, 1991; Marchman, 1993; Schwartz et al., 1990; Seidenberg and

²A number of readable introductions to connectionism are now available. See Bechtel and Abrahamsen, 1991; Churchland and Sejnowski, 1992; Dayhoff, 1990. An excellent but more technical introduction can be found in Hertz et al., 1991.

McClelland, 1989). Second, the same units can participate in many different patterns, and many different patterns coexist in a superimposed fashion across the same set of units. This fact can be used to account for degrees of similarity between patterns, and for the ways in which patterns penetrate, facilitate and/or interfere with one another at various points in learning and development (for an expanded discussion of this point, see Bates et al., 1991; Elman et al., 1996).

(2) *Graded Rules.* Contrary to rumor, it is not the case that connectionist systems have no rules. However, the rules or “mappings” employed by connectionist nets take a very different form from the crisp algorithms contained within the programs employed by a serial digital computer. These include the learning rule itself (i.e. the principle by which the system reduces error and “decides” when it has reached a good fit between input and output), and the functions that determine when and how a unit will fire. But above all, the “rules” in a connectionist net include the connections that hold among units, i.e. the links or “weights” that embody all the potential mappings from input to output across the system as a whole. This means that rules (like representations) can exist by degree, and vary in strength.

It should also be clear from this description that it is difficult to distinguish between rules and representations in a connectionist net. The knowledge or “mapping potential” of a network comprises the units that participate in distributed patterns, and the connections among those units. Because all these potential mappings coexist across the same “territory”, they must compete with one another to resolve a given input. In the course of this competition, the system does not “decide” between alternatives in the usual sense; rather, it “relaxes” or “resolves” into a (temporary) state of equilibrium. In a stochastic system of this kind, it is possible for several different networks to reach the same solution to a problem, each with a totally different set of weights. This fact runs directly counter to the tendency in traditional cognitive and linguistic research to seek “the rule” or “the grammar” that underlies a set of behavioral regularities. In other words, rules are not absolute in any sense—they can vary by degree within a given individual, and they can also vary in their internal structure from one individual to another. We believe that these properties are far more compatible with the combination of universal tendencies and individual variation that we see in the course of human development, and they are compatible with the remarkable neural and behavioral plasticity that is evident in children who have suffered early brain injury (Bates et al., 1997; Elman et al., 1996; Marchman, 1993).

(3) *Learning as structural change.* As we pointed out earlier, much of the current excitement about connectionist systems revolves around their capacity for learning and self-organization. Indeed, the current boom in connectionism has brought learning and development back onto center stage in cognitive science. These systems really do change as a function of learning, dis-

playing forms of organization that were not placed there by the programmer (or by nature, or by the Hand of God). To be sure, the final product is co-determined by the initial structure of the system and the data to which it is exposed. These systems are not anarchists, nor solipsists. But in no sense is the final product “copied” or programmed in. Furthermore, once the system has learned, it is difficult for it to “unlearn”, if by “unlearning” we mean a return to its pristine prelearning state. This is true for the reasons described in (1) and (2): the knowledge contained in connectionist nets is contained in and defined by its very architecture, in the connection weights that currently hold among all units as a function of prior learning. Knowledge is not “retrieved” from some passive store, nor is it “placed in” or “passed between” spatially localized buffers. Learning is structural change, and experience involves the activation of potential states in that system as it is currently structured.

From this point of view, the term “acquisition” is an infelicitous way of talking about learning or change. Certain states become possible in the system, but they are not acquired in the usual sense, i.e. found or purchased or stored away like nuts in preparation for the winter. This property of connectionist systems permits us to do away with problems that have been rampant in certain areas of developmental psychology, e.g. the problem of determining “when” a given piece of knowledge is acquired, or “when” a rule finally becomes productive. Instead, development (like the representations and mappings on which it is based) can be viewed as a gradual process; there is no single moment at which learning can be said to occur (but see non-linearity, below).

(4) *Software as hardware.* We have stated that knowledge in connectionist nets is defined by the very structure of the system. For this reason, the hardware/software distinction is impossible to maintain under the Second Computer Metaphor. This is true whether or not the structure of connectionist nets as currently conceived is “neurally real”, i.e. like the structure that holds in real neural systems. We may still have the details wrong (indeed, we probably do), but the important point for present purposes is that there is no further excuse for ignoring potential neural constraints on proposed cognitive architectures. The distinction that has separated cognitive science and neuroscience for so long has fallen, like the Berlin Wall. Some cognitive psychologists and philosophers of science believe that is not a good thing (and indeed, the same might be said someday for the Berlin Wall). But we are convinced that this historic event is a good one, especially for those of us who are interested in the co-development of mind and brain. We are going the right direction, even though we have a long way to go.

(5) *Nonlinear dynamics.* Connectionist networks are nonlinear dynamical systems, a fact that follows from several properties of connectionist architecture including the existence of intervening layers between

inputs and outputs (permitting the system to go beyond linear mappings), the nonlinear threshold functions that determine how and when a single unit will fire, and the learning rules that bring about a change in the weighted connections between units. Because these networks are nonlinear systems, they can behave in unexpected ways, mimicking the U-shaped learning functions and sudden moments of “insight” that challenged old Stimulus-Response theories of learning, and helped to bring about the cognitive revolution in the 1960s (MacWhinney and Leinbach 1991; Plunkett and Marchman, 1991a, b).

(6) *Emergent form.* Because connectionist networks are nonlinear systems, capable of unexpected forms of change, they are also capable of producing truly novel outputs (see also Elman et al., 1996; MacWhinney, 1999). In trying to achieve stability across a large number of superimposed, distributed patterns, the network may hit on a solution that was “hidden” in bits and pieces of the data; that solution may be transformed and generalized across the system as a whole, resulting in what must be viewed as a qualitative shift. This is the first precise, formal embodiment of the notion of emergent form—an idea that stood at the heart of Piaget’s theory of change in cognitive systems. As such, connectionist systems may have the very property that we need to free ourselves from the nature–nurture controversy. New structures can emerge at the interface between “nature” (the initial architecture of the system) and “nurture” (the input to which that system is exposed). These new structures are not the result of black magic, or vital forces. They are the result of laws that govern the integration of information in nonlinear systems—which brings us to our final section.

4. Some Common Misconceptions about Connectionism

It is no doubt quite clear to the reader that we are enthusiastic about the Second Computer Metaphor, because we believe that it will help us to pick up a cold trail that Piaget first pioneered, moving toward a truly interactive theory of change. But we are aware of how much there is to do, and how many pitfalls lie before us. We are also aware of some of the doubts and worries about this movement that are currently in circulation. Perhaps it would be useful to end this essay with some answers to some common misconceptions about connectionism, with special reference to the application of connectionist principles within developmental psychology.

Worry 1. “Connectionism is nothing but associationism, and we already know the limits of associationism” (e.g. Fodor and Pylyshyn, 1988). As we pointed out above, multilayer connectionist nets are nonlinear dynamical systems, whereas the familiar associationist models of the past rested on assumptions of linearity. This is both the good news, and the bad news. The good news is that nonlinear systems can learn relationships of considerable complexity, and they can produce surprising and (of course) nonlinear forms of

change. The bad news is that no one really understands the limits and capabilities of nonlinear dynamical systems. Maybe this is also good news: we have finally met our goal, after years of physics envy, because we have finally reached the same frontiers of ignorance as the physicists! Presumably, the limits of these systems will someday be known (although probably not within our lifetimes). But right now, it would be grossly premature to claim that connectionist networks can “never” perform certain functions. Anyone who claims that we already know the limits of this kind of associationism has been misinformed.

Worry 2. “There are no interesting internal representations in connectionist nets” (e.g. Pinker and Prince, 1988). There are indeed complex and rich representations in connectionist networks, and transformations that do the same work as rules in classical systems. However, these rules and representations take a radically different form from the familiar symbols and algorithms of serial digital computers and/or generative linguistics. The representations and rules embodied in connectionist nets are implicit and highly distributed. Part of the challenge of modern research on neural networks is to understand exactly what a net has learned after it has reached some criterion of performance. So far, the answer appears to be that the solutions that networks find very often look quite different than their symbolic counterparts (for examples, see Elman, 1989, 1990, 1991).

Worry 3. “Connectionist nets only yield interesting performance on cognitive problems when the experimenter ‘sneaks in’ the solution by (a) fixing the internal weights until they work, or (b) laying out the solution in the input” (e.g. Lachter and Bever, 1988). Part of the fascination of connectionist modeling lies in the fact that it offers the experimenter so many surprises. These are self-organizing systems that learn how to solve a problem. As the art is currently practiced, *no one* fiddles with the internal weights but the system itself, in the course of learning. Indeed, in a simulation of any interesting level of complexity, it would be virtually impossible to reach a solution by “hand-tweaking” of the weights. As for the issue of “sneaking the solution into the input”, we have seen several simulations in which the Experimenter did indeed try to make the input as explicit as possible—and yet the system stubbornly found a different way to solve the problem. Good connectionist modelers approach their simulations with the same spirit of discovery and breathless anticipation that is very familiar to those who carry out real experiments with real children. Aside from being close to impossible, cheating would not be any fun at all—and the hand-crafting of solutions is usually considered a form of cheating.

Worry 4. “The supposed commitment to neural plausibility is a scam; no one really takes it seriously.” Connectionists work at many different levels between brain and behavior. In current simulations of higher cognitive processes, it is true that the architecture is “brain-like” only in a very indirect sense. In fact, the

typical 100-neuron connectionist toy is “brain-like” only in comparison with the serial digital computer (which is wildly unlike nervous systems of any known kind). The many qualities that separate real brains from connectionist simulations have been described in detail elsewhere (Churchland and Sejnowski, 1992; Crick, 1989; Hertz, Krogh and Palmer, 1991). The real questions are: (a) is there anything of interest that can be learned from simulations in simplified systems, and (b) can connectionists “add in” constraints from real neural systems in a series of systematic steps, approaching something like a realistic theory of mind and brain? Of course we still do not know the answer to either of these questions, but there are many researchers in the connectionist movement who are trying to bring these systems closer to neural reality. For example, efforts are underway to study the computational properties of different neuronal types. Some researchers are exploring analogues to synaptogenesis and synaptic pruning in neural nets. Others are looking into the computational analogues of neural transmitters within a fixed network structure. The current hope is that work at all these different levels will prove to be compatible, and that a unified theory of the mind and brain will someday emerge. Of course we are a long way off, but the commitment by most of the researchers that we know in this field is a very serious one. It has launched a new spirit of interdisciplinary research in cognitive neuroscience, one with important implications for developmental psychology.

Worry 5. “Connectionism is anti-nativist, and efforts are underway to reinstate a *tabula rasa* approach to mind and development” (e.g. Kirsh, 1992; Pinker, 1999). It is true that many current simulations assume something like a *tabula rasa* in the first stages of learning (e.g. a random “seeding” of weights among fully connected units before learning begins). This has proven to be a useful simplifying assumption, in order to learn something about the amount and type of structure that has to be assumed for a given type of learning to go through. But there is no logical incompatibility between connectionism and nativism. Indeed, just as many historians have argued that Franklin Delano Roosevelt saved capitalism, connectionism may prove to be the salvation of nativist approaches to mind. The problem with current nativist theories is that they offer no serious account of what it might mean in biological terms for a given structure or idea to be innate. In neural networks, it is possible to explore various avenues for building in innate structure, including minor biases that have major structural consequences across a range of environmental conditions (Elman, 1994; Elman et al., 1996; Jacobs, 1999; Jacobs et al., 1991; Jacobs and Jordan, 1992). In fact, within connectionist models there are coherent ways to talk about 90% or 10% of any innate idea! This is an approach that has not been explored in any detail to date, but the possibilities are intriguing, and might (ironically enough) end up being connectionism’s greatest contribution to developmental cognitive neuroscience.

5. Epilogue: New directions in Developmental Connectionism

The story is far from over—so this is not a true epilogue—but since we first wrote this chapter it has become even clearer how connectionist models are offering new insights into developmental phenomena. Consider the following (by no means exhaustive) examples:

* *U-shaped curves.* In the course of learning a number of tasks, children frequently exhibit various “U-shaped” patterns of behavior in which good initial performance is followed by poorer performance, which eventually again improves. One of the best-known cases of this is the acquisition of the past tense of English. Networks that are trained on similar tasks exhibit the same patterns of behavior (MacWhinney et al., 1989; Plunkett and Juola, 1999; Plunkett and Marchman, 1991a, 1991b, 1993; Rumelhart and McClelland, 1986). Furthermore, the sorts of differences in learning of regular forms (e.g., those which add “-ed”) compared with irregulars (“sing”–“sang”) have now been explored in network models of many other languages (e.g., German, Arabic; Hare, Elman and Daugherty, 1995; Nakisa, Plunkett and Hahn, in press; Plunkett and Nakisa, 1997).

* *Critical periods in brain plasticity.* Developmental changes in brain plasticity—the ability to recover from early lesions—suggest that there are “critical periods” during which the brain is more plastic and able to learn. Marchman (1993) has shown that such changes in plasticity can also arise merely as a by-product of learning itself, and do not necessarily require some kind of exogenous (learning-independent) change in the structure or function of the brain. The same changes in brain structure that occur as a result of learning may also make recovery from stroke more difficult, and make later learning harder.

* *Lesioning networks.* Networks that are trained on tasks such as reading or verb morphology demonstrate, when “lesioned,” symptoms and patterns of recovery which closely resemble the patterns of human aphasics (Farah and McClelland, 1991; Hinton and Shallice, 1991; Marchman, 1993; Martin et al., 1994; Plaut and Shallice, 1993; Seidenberg and McClelland, 1989). Finally, the kinds of dissociations in behavior that may be found as a result of brain damage, in which regular forms are better preserved than irregulars (or vice versa) have been simulated by lesioning networks in which the regular/irregular distinction is not directly encoded (Juola and Plunkett, 1998; Joanisse and Seidenberg, 1999).

* *Learning rediscovered.* There have been a number of recent demonstrations of apparently precocious knowledge in young infants (e.g., very early learning of language-like patterns; object permanence). Some researchers have taken this as evidence for innate knowledge, and the existence of abstract rules and symbols. But similar results have been easily replicated in networks that are sensitive to statistical properties in

their environment, suggesting that the environment is a richer source of information than might be thought, and that young infants are using a relatively simple learning algorithm to extract this information (Mareschal, Plunkett and Harris, 1995; Munakata, McClelland, Johnson and Siegler, 1999; Seidenberg and Elman, 1999).

* *Readiness phenomena.* Children are known to go through phases in which behavior changes slowly and is resistant to new learning. At other points in time children show heightened sensitivity to examples and rapid changes in behavior. Networks exhibit similar readiness phenomena (Elman et al., 1996, Chapter 6; McClelland, 1989).

* *The importance of starting small.* Some aspects of grammatical structure are particularly hard to learn, and it has been claimed that these structures must therefore be innate. Paradoxically, these structures can be learned in networks as long as learning begins when the network has a limited working memory (similar to that of an infant or young child). This limitation turns out to be crucial for learning the more complex aspects of language.

* *Activity-dependent origins of brain structure.* Visual cortex in mammals is well known to include neurons that are selectively sensitive to highly specific visual inputs. These neurons include edge detectors, center-surround cells, and motion detectors. Biologically plausible network models have been constructed which demonstrate that such specialized response properties do not have to be prespecified. They emerge naturally and inevitably from cells which are initially uncommitted, simply as a function of a simple learning rule and exposure to stimulation (Linsker, 1986, 1990; Miller, Keller and Stryker, 1989; Sereno and Sereno, 1991). These artificial networks even develop the characteristic zebra-like striped patterns seen in ocular dominance columns in real cortex (Miller, Keller and Stryker, 1989).

* *Learning the categories of language.* Must children have innate knowledge of linguistic categories such as “noun” or “verb”? Or can these notions be learned? Network simulations as well as corpus analyses have shown that such categories can be learned merely from the distributional properties in the language that young children hear (Elman, 1990; Mintz, Newport and Bever, submitted).

* * *

These examples (as well as many others described in Elman et al., 1996) lead us to believe we will soon see a revival of Piagetian theory within a connectionist framework—not a mindless reinterpretation of the old theory in modern jargon, but a return to Piaget’s program of genetic epistemology, instantiating his principles of equilibration and adaptation in concrete systems that really work—and really change. As we said before, Piaget spent the later decades of his life seeking a way of formalizing the theory, to answer critics (including Piaget himself) who charged that his principles of change were much too vague. We think that Piaget would have loved these new possibilities if he had lived to see them. We now have an opportunity to pick up

the threads of his old program and move it forward into an exciting new decade, incorporating all the new insights and new empirical information that has been gained in the interim, without abandoning the fundamental commitment of developmental psychology to the study of change.

REFERENCES

- Anderson, J.A. (1972) A simple neural network generating an interactive memory. *Mathematical Bio-Sciences*, 8, 137–160.
- Anderson, J.A. and Rosenfeld, E. (1989) *Neurocomputing: Foundations of research*. MIT Press/Bradford Books.
- Baillargeon, R. (1999) Young infants' expectation about hidden objects: A reply to three challenges. *Developmental Science*, 2, 115–132.
- Baillargeon, R. and de Vos, J. (1991) Object permanence in young infants: Further evidence. *Child Development*, 62, 1227–1246.
- Bates, E., Elman, J., Johnson, M., Karmiloff-Smith, A., Parisi, D. and Plunkett, K. (1998) Innateness and emergentism. In W. Bechtel and G. Graham (eds), *A companion to cognitive science*. Malden, MA & Oxford: Blackwell Publishers, pp. 590–601.
- Bates, E., Thal, D., Finlay, B.L. and Clancy, B. (in press) Early language development and its neural correlates. To appear in F. Boller & J. Grafman (Series eds) & I. Rapin & S. Segalowitz (Vol. eds), *Handbook of neuropsychology, Vol. 7: Child neurology* (2nd ed.). Amsterdam: Elsevier.
- Bates, E., Thal, D. and Marchman, V. (1991) Symbols and syntax: A Darwinian approach to language development. In N. Krasnegor, D. Rumbaugh, E. Schiefelbusch and M. Studdert-Kennedy (eds), *The biological and behavioral determinants of language development*. Hillsdale, NJ: Erlbaum.
- Bates, E., Thal, D., Trauner, D., Fenson, J., Aram, D., Eisele, J. and Nass, R. (1997) From first words to grammar in children with focal brain injury. In D. Thal and J. Reilly (eds), Special issue on Origins of Communication Disorders, *Developmental Neuropsychology*, 13, 275–343.
- Bechtel, W. and Abrahamsen, A. (1991) *Connectionism and the mind*. Oxford: Basic Blackwood.
- Bruner, J. and Sherwood, V. (1976) Peekaboo and the learning of rule structures. In J.S. Bruner, A. Jolly and K. Sylva (eds), *Play: Its role in development and evolution*. New York: Basic Books, Inc.
- Churchland, P. and Sejnowsky, T. (1992) *The net effect*. Cambridge, MA: MIT Press/Bradford Books.
- Clancy, B., Darlington, R.B. and Finlay, B.L. (2000) The course of human events: Predicting the timing of primate neural development. *Developmental Science*, 3, 57–66.
- Crick, F. (1989) The recent excitement about neural networks. *Nature*, 337, 129–132.

- Dayhoff, J. (1990) *Neural network architectures*. New York: Van Nostrand Reinhold.
- Eccles, J.L. (1953) *The neurophysiological basis of mind*. Oxford: Clarendon.
- Elman, J. (1989) Structured representations and connectionist models. In *The Eleventh Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Elman, J. (1990) Finding structure in time. *Cognitive Science*, 14, 179–211.
- Elman, J. (1991) Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7, 195–225.
- Elman, J. (1993) Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71–99.
- Elman, J. (1994) Learning and development in neural networks: The importance of starting small. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing*. Hillsdale, NJ: Erlbaum.
- Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D. and Plunkett, K. (1996) *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press/Bradford Books.
- Farah, M.J. and McClelland, J. (1991) A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General*, 120, 339–357.
- Feldman, J.A. and Ballard, D.H. (1980) *Computing with connections*. TR 72. University of Rochester: Computer Science Department.
- Ferguson, C. and Snow, C. (1978) *Talking to children*. Cambridge: Cambridge University Press.
- Fodor, J.A. (1981) *Representations*. Brighton (Sussex): Harvester Press.
- Fodor, J.A. and Pylyshyn, Z.W. (1988) Connectionism and cognitive architecture: A critical analysis. In S. Pinker and J. Mehler (eds), *Connections and symbols*. Cambridge, MA: MIT Press/Bradford Books, pp. 3–71.
- Gerhart, J., & Kirschner, M. (1997) Cells, embryos and evolution: Toward a cellular and developmental understanding of phenotypic variation and evolutionary adaptability. Malden, MA: Blackwell Science.
- Givón, T. (1984) *Syntax: A functional-typological introduction. Volume I*. Amsterdam: John Benjamins.
- Golden, R. (1996) *Mathematical models for neural network analysis and design*. Cambridge, MA: MIT Press.
- Grossberg, S. (1968) Some physiological and biochemical consequences of psychological postulates. *Proceedings of the National Academy of Science, USA*, 60, 758–765.
- Grossberg, S. (1972) Neural expectation: Cerebellar and retinal analogs of cells fired by learnable or unlearned pattern classes. *Kybernetik*, 10, 49–57.
- Grossberg, S. (1987) *The adaptive brain*, 2 vols. Amsterdam: Elsevier.
- Hare, M., Elman, J. and Daugherty, K.G. (1995) Default generalization in connectionist networks. *Language and Cognitive Processes*, 10, 601–630.
- Hebb, D. (1949) *The organization of behavior*. New York: Wiley.
- Hertz, J., Krogh, A. and Palmer, R. (1991) *Introduction to the theory of neural computation*. Redwood City, California: Addison Wesley.
- Hinton, G.E. and Shallice, T. (1991) Lesioning a connectionist network: Investigations of acquired dyslexia. *Psychological Review*, 98, 74–95.
- Hinton, G.E. and Anderson, J.A. (1981) *Parallel models of associative memory*. Hillsdale, NJ: Erlbaum.
- Hornik, K., Stinchcombe, M. and White, H. (1989) Multilayer feedforward networks are universal approximators. *Neural Networks*, 2, 359–366.
- Hyams, N. (1986) *Language acquisition and the theory of parameters*. Dordrecht & Boston: Reidel.
- Jacobs, R. (1999) Computational studies of the development of functionally specialized neural models. *Trends in Cognitive Sciences*, 3, 31–38.
- Jacobs, R., Jordan, M. and Barto, A. (1991) Task decomposition through competition in a modular connectionist architecture: The what and where visual tasks. *Cognitive Science*, 15, 219–250.
- Jacobs, R. and Jordan, M. (1992) Computational consequences of a bias toward short connections. *Journal of Cognitive Neuroscience*, 4, 323–336.
- Joanisse, M.F. and Seidenberg, M.S. (1999) Impairments in verb morphology after brain injury: A connectionist model. *Proceedings of the National Academy of Sciences of the United States of America*, 96, 7592–7597.
- Johnson, M.H. (1997) *Developmental cognitive neuroscience: An introduction*. Cambridge, MA: Blackwell Publishers.
- Juola, P. and Plunkett, K. (1998) Why double dissociations don't mean much. In M.A. Gernsbacher and S.J. Derry (eds), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 561–566). Mahwah, NJ: Erlbaum.
- Kellman, P.J., Spelke, E.S. and Short, K.R. (1986) Infant perception of object unity from translatory motion in depth and vertical translation. *Child Development*, 57, 72–86.
- Kirsh, D. (1992) PDP Learnability and innate knowledge of language. *Center for Research in Language Newsletter Vol. 6, no. 3*. University of California, San Diego.
- Kohonen, T. (1977) *Associative memory: A system-theoretical approach*. Berlin: Springer.
- Lachter, J. and Bever, T.G. (1988) The relation between linguistic structure and associative theories of

- language learning: A constructive critique of some connectionist learning models. In S. Pinker and J. Mehler (eds), *Connections and symbols*. Cambridge, MA: MIT Press/Bradford Books, pp. 3–71.
- Lakoff, G. (1987) *Fire, women, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
- Langacker, R. (1987) *Foundations of cognitive grammar: Theoretical perspectives. Volume I*. Stanford: Stanford University Press.
- Le Cun, Y. (1985) Une procédure d'apprentissage pour réseau à seuil asymétrique. In *Cognitiva 85: à la Frontière de l'Intelligence Artificielle des Sciences de la Connaissance des Neurosciences*, Paris, pp. 599–604.
- Lightfoot, D. (1991) The child's trigger experience — Degree-0 learnability. *Behavioral and Brain Sciences*, 14, 321–375.
- Linsker, R. (1986) From basic network principles to neural architecture (series). *Proceedings of the National Academy of Sciences, USA*, 83, 7508–7512, 8390–8394, 8779–8783.
- Linsker, R. (1990) Perceptual neural organization: Some approaches based on network models and information theory. *Annual Review of Neuroscience*, 13, 257–281.
- MacWhinney, B. (1991) Implementations are not conceptualizations: Revising the verb-learning model. *Cognition*, 40, 121–157.
- MacWhinney, B. (ed) (1999) *The emergence of language*. Mahwah, NJ: Lawrence Erlbaum.
- MacWhinney, B., Leinbach, J. (1991) Implementations are not conceptualizations: Revising the verb-learning model. *Cognition*, 40, 121–157.
- MacWhinney, B., Leinbach, J., Taraban, R. and McDonald, J. (1989) Language learning: Cues or rules? *Journal of Memory and Language*, 28, 255–277.
- Mareschal, D., Plunkett, K. and Harris, P. (1995) Developing object permanence: A connectionist model. In J.D. Moore and J.F. Lehman (eds), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 170–175). Mahwah, NJ: Erlbaum.
- Marchman, V. (1993) Constraints on plasticity in a connectionist model of the English past tense. *Journal of Cognitive Neuroscience*, 5, 215–234.
- Martin, N., Dell, G.S., Saffran, E.M., & Schwartz, M.F. (1994) Origins of paraphasias in deep dysphasia: Testing the consequences of a decay impairment to an interactive spreading activation model of lexical retrieval. *Brain and Language*, 47, 52–88.
- McClelland, J.L. (1989) Parallel distributed processing: Implications for cognition and development. In R.G.M. Morris (ed), *Parallel distributed processing: Implications for psychology and neurobiology*. Oxford: Clarendon Press, pp. 9–45.
- McClelland, J. and Rumelhart, D. (1986) *Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 2*. Cambridge, MA: MIT Press/Bradford Books.
- McCulloch, W. and Pitts, W. (1943) A logical calculus of ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115–133. Reprinted in J. Anderson and E. Rosenfeld (eds), *Neurocomputing: Foundations of research*. Cambridge, MA: MIT Press.
- Mead, C. (1989) Analog VLSI and neural systems. Inaugural address presented to the Institute for Neural Computation, October, 1989. University of California, San Diego.
- Miller, K.D., Keller, J.B. and Stryker, M.P. (1989) Ocular dominance column development: Analysis and simulation. *Science*, 245, 605–615.
- Minsky, M. (1956) Some universal elements for finite automata. In C.E. Shannon and J. McCarthy (eds), *Automata studies*. Princeton: Princeton University Press, pp. 117–128.
- Minsky, M. and Papert, S. (1969) *Perceptrons*. Cambridge, MA: MIT Press.
- Mintz, T.H., Newport, E.L. and Bever, T.G. (submitted) The distributional structure of grammatical categories in speech to young children.
- Munakata, Y., McClelland, J.L., Johnson, M.H. and Siegler, R.S. (1999) Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, 4, 686–713.
- Nakisa, R., Plunkett and Hahn, U. (in press) A cross-linguistic comparison of single and dual-route models of inflectional morphology. In P. Broeder & J. Murre (eds), *Cognitive models of language acquisition*. Cambridge, MA: MIT Press.
- Papert, S. (1988) One AI or Many? *Daedalus: Artificial Intelligence*. Winter.
- Piaget, J. (1952) *The origins of intelligence in children*. New York: International Universities Press.
- Piaget, J. (1970a) *Structuralism*. New York: Basic Books.
- Piaget, J. (1970b) *Genetic epistemology*. New York: Columbia University Press.
- Piaget, J. (1971) *Biology and knowledge: An essay on the relations between organic regulations and cognitive processes*. Chicago: University of Chicago Press.
- Piatelli-Palmarini, M. (1989) Evolution, selection, and cognition: From “learning” to parameter setting in biology and the study of language. *Cognition*, 31, 1–44.
- Pinker, S. (1999) *Words and rules: The ingredients of language*. New York: Basic Books.
- Pinker, S. and Prince, A. (1988) On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. In S. Pinker and J. Mehler (eds), *Connections and sym-*

- bols.* Cambridge, MA: MIT Press/Bradford Books, pp. 3–71.
- Plaut, D.C., and Shallice, T. (1993) Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, 10, 377–500.
- Plunkett, K. and Elman, J. (1997) *Exercises in rethinking innateness: A handbook for connectionist simulations.* Cambridge, MA: MIT Press/Bradford Books.
- Plunkett, K. and Juola, P. (1999) A connectionist model of English past tense and plural morphology. *Cognitive Science*, 23, 463–490.
- Plunkett, K. and Marchman, V. (1991a) U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition*, 38, 43–102.
- Plunkett, K. and Marchman, V. (1991b) *From rote learning to system building* (Tech. Rep. 9020). Center for Research in Language, University of California, San Diego.
- Plunkett, K. and Nakisa, R.C. (1997) A connectionist model of the Arabic plural system. *Language and Cognitive Processes*, 12, 807–836.
- Roeper, T. and Williams, E., (eds) (1987) *Parameter setting.* Dordrecht and Boston: Reidel.
- Rogoff, B. (1990) *Apprenticeship in thinking: Cognitive development in social context.* New York: Oxford University Press.
- Rosenblatt, F. (1958) The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, 386–408.
- Rosenblatt, F. (1962). *Principles of neurodynamics.* New York: Spartan.
- Rumelhart, D., Hinton, G. and Williams, R. (1986a) Learning representations by back-propagating errors. *Nature*, 323, 533–536.
- Rumelhart, D., McClelland, J. and the PDP Research Group (1986b) *Parallel distributed processing: explorations in the microstructure of cognition, Vol. 1.* Cambridge, MA: MIT/ Bradford Books.
- Schwartz, M.F., Saffran, E.M. and Dell, G.S. (1990) *Comparing speech error patterns in normals and jargon aphasics: Methodological issues and theoretical implications.* Presented to the Academy of Aphasia, Baltimore, MD.
- Seidenberg, M. and Elman, J.L. (1999) Do infants learn grammar with algebra or statistics? *Science*, 284, 434–435.
- Seidenberg, M. and McClelland, J.L. (1989) A distributed developmental model of visual word recognition and naming. *Psychological Review*, 96, 523–568.
- Selfridge, O.G. (1958) Pandemonium: A paradigm for learning. In *Mechanisation of Thought Processes: Proceedings of a Symposium Held at the National Physical Laboratory, November 1958.* London: HMSO, pp. 513–526.
- Sereno, M.I. and Sereno, M.E. (1991) Learning to see rotation and dilation with a Hebb rule. In R.P. Lippman, J. Moody, & D.S. Touretzky (eds), *Advances in neural information-processing systems 3.* San Mateo, CA: Morgan Kaufman, pp. 320–326.
- Spelke, E. (1990) Principles of object perception. *Cognitive Science*, 14, 29–56.
- Spelke, E. (1991) Physical knowledge in infancy: Reflections on Piaget’s theory. In S. Carey and R. Gelman (eds), *The epigenesis of mind: Essays on biology and cognition.* Hillsdale, NJ: Erlbaum, pp. 133–69.
- Spelke, E.S. and Newport, E.L. (1998) Nativism, empiricism, and the development of knowledge. In W. Damon (Series ed.) and D. Kuhn and R. Siegler (Vol. eds), *Handbook of child psychology: Vol. 1, 5th ed. Theoretical models of human development.* New York: Wiley, pp. 275–340.
- Thal, D., Marchman, V., Stiles, J., Aram, D., Trauner, D., Nass, R. and Bates, E. (1991) Early lexical development in children with focal brain injury. *Brain and Language*, 40, 491–527.
- Von Neumann, J. (1951) The general and logical theory of automata. In L.A. Jeffress (ed), *Cerebral mechanisms in behavior.* New York: Wiley.
- Von Neumann, J. (1958) *The computer and the brain.* New Haven: Yale University Press.
- Werner, H. (1948) *Comparative psychology of mental development.* New York: International Universities Press.
- Willshaw, D.J., Buneman, O.P. and Longuet-Higgins, H.C. (1969) Nonholographic associative memory. *Nature*, 222, 960–962.
- Zipser, D. and Andersen, R.A. (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331, 679–684.




Figure 1. A simple perceptron network




Figure 2. A multi-layer network. The “hidden units” allow such networks to form internal representations. Multi-layer networks of this sort have been shown to be able to approximate any function to any arbitrary degree of accuracy.