The Journal of Mind and Behavior
Winter and Spring 2000, Volume 21, Numbers 1 and 2
Pages 153–184
ISSN 0271–0137
ISBN 0–930195–11–6

The Nature of Distributed Learning and Remembering

Asghar Iran-Nejad
University of Alabama

Abdollah Homaifar
North Carolina A&T State University, Greensboro

Researchers have held different views on what role the nervous system should play in the study of psychological phenomena. By far, the most informative line of research in the area has been conducted by Lashley whose work has opened our eyes to the possibility that learning and remembering are unexplainable in terms of the storage and retrieval of specific traces. However, with this exception, the twentieth century is likely to be remembered as an era during which the brain has been considered irrelevant for the study of the mind. This has certainly been the case with the research following the computer-inspired cognitive revolution. Perhaps the most revealing indication of the degree of reluctance to embrace the brain in the study of the mind can be found in the so-called brain-inspired connectionism that purports to use the brain as a metaphor, and not as the literal foundation it really is, for the structure of cognition. Focusing on the topics of learning and remembering, this paper discusses the role of the brain in the research of Lashley, brain-inspired connectionism, and the emerging field of biofunctional cognition. The hope is to illustrate, through biofunctional cognition, the productive nature of basing psychological thinking on the foundation of a comprehensive theory of the functioning of the nervous system.

Lashley (1915, 1929, 1950) devoted more than three decades of research in pursuit of localized memory traces in the brain. His detailed investigations uncovered no such traces, but prepared the empirical groundwork for the development of the nonlocalizationist perspective on learning and remembering. After Lashley (1890–1958), the research on distributed memory con-

---

An earlier version of this article appeared in a German chapter on the same topic published in 1991 in S.J. Schmidt (Ed.), *Gedächtnis* [*Memory*] (pp. 206–249), Frankfurt/Main, West Germany: ©Suhrkamp Verlag. Requests for reprints should be sent to Asghar Iran-Nejad, Ph.D., Educational Psychology Program, College of Education, University of Alabama, Box 870231, Tuscaloosa, Alabama 35487–0231. Email: airannej@bamaed.ua.edu

tinued a steady course (see Iran-Nejad and Ortony, 1984). Yet, the mainstream cognitive science ignored this research and was involved in a cognitive revolution that reified local storage metaphors in their most concrete form ever in terms of the computer software analogy. Then in mid 1980s, there was a sudden explosion of interest in distributed representations.

Today, the notion of distributed learning and remembering is firmly entrenched in a new form of associationism called *parallel distributed processing* (PDP) connectionism. Whereas there is widespread agreement that memory is represented in a distributed fashion, much less consensus exists as to whether PDP connectionism is the best way to think about distributed learning and remembering (DLR). This is why a good portion of this article is devoted to clarifying the status of DLR in PDP associationism. The main theme here is that DLR research originated and continued to grow in a nonassociative context and is fundamentally incompatible with connectionism (see Iran-Nejad, 1980; Iran-Nejad and Ortony, 1984). First, we briefly discuss some of the reasons why PDP associationism and the (nonassociative) foundational research that led to the discovery of the notion of DLR are paradigmatically incompatible in the Kuhnian (1962) sense of the term. We then argue that PDP connectionism and conventional cognitive science of the 1970s are fundamentally the same, despite the differences in appearance and rhetoric. Next, a discussion of some of the major aspects of DLR follows. Finally, we conclude the article with a biofunctional analysis of distributed learning and remembering.

PDP Associationism and DLR History: Some Issues

Real Discoveries Must Count

To get a flavor for the degree of incompatibility between the associative and nonassociative approaches to DLR, consider the seminal work of Lashley (1929, 1950, 1951). PDP authors have seldom discussed the implications of Lashley's data for their models. For example, when they first introduced their distributed model, McClelland and Rumelhart (1985a) never mentioned Lashley in the section of their paper where "some important credits [were] in order" (p. 161). Interestingly, they did mention him in the section that discussed their superposition hypothesis, meaning that the same brain regions store memory traces for many experiences in the form of superposed layers of information (numeric connection weights for PDP models). However, as discussed elsewhere (Iran-Nejad and Ortony, 1984), the fact that the superposition hypothesis is suggested in some of Lashley's writings is perhaps the weakest aspect of his notion of distributed learning and remembering.

Lashley's research has received the same sort of treatment throughout the entire booming PDP literature. The book, *Parallel Models of Associative Memory* (Anderson and Hinton, 1981), that launched the current enthusiasm in PDP connectionism contained three citations of Lashley's work, only one of which touched upon substance, finding "uncompromising" Lashley's statement "that even the reservation of individual synapses for special associative reactions is impossible" (Lashley, 1950, p. 480). Anderson and Hinton did not mention that three decades of systematic research had convinced Lashley to accept the hypothesis of equipotentiality of the so-called association tracts in the nervous system. Nor did they discuss the implications of this hypothesis, if it were to be correct, for what PDP modelers often refer to as the new "insight that the knowledge is stored in the interconnections between units" (McClelland, Rumelhart, and Hinton, 1986, p. 33).

Similarly, in the twenty-six chapters of the two PDP volumes (McClelland, Rumelhart, and the PDP Research Group, 1986; Rumelhart, McClelland, and the PDP Research Group, 1986), there was only one reference to Lashley acknowledging his contribution to distributed representations, while also stating that he "may have been too radical and too vague, and his doctrine of equipotentiality of broad regions of cortex clearly overstated the problem" (McClelland, Rumelhart, and Hinton, 1986, p. 41). Again, nowhere in the 1158 pages of two PDP volumes was there a discussion of how Lashley was vague and what he meant by the term *equipotential*.

This dwelling on the inadequate treatment of Lashley by PDP modelers may strike the reader as overkill. However, we felt that this discussion was needed for two reasons. First, we need to stress that the light treatment of Lashley in the PDP literature is evidence for the incompatibility of the two perspectives. Second, Lashley advanced the neuroscience of DLR far beyond where we are today while PDP, essentially a return to classic associationism, seems to have pushed it back. Moreover, progress in understanding how the brain creates the mind requires a principled method of building the present on the foundation of the past.

Oil and Water Do Not Mix: Nonassociative Concepts in the PDP Literature

Jenkins (1974) rejected associationism as a model of memory and discussed contextualism as an alternative. For him, the two perspectives were incompatible. Like Lashley, Jenkins came to this conclusion after many years of experience with associationism. However, in their interactive activation model, McClelland and Rumelhart (1981; Rumelhart and McClelland, 1982) equated context with activation of discrete nodes and connections in an associative network.

In another programmatic line of research, Medin and his colleagues demonstrated that associationism, in its various forms, cannot handle contextual and theoretical knowledge (Medin and Schaffer, 1978; Medin and Shwanenflugel, 1981; Murphy and Medin, 1985). McClelland and Rumelhart (1985a) referred to this literature and acknowledged that context presents a challenge for their distributed model, concluding that "the problem is a severe one, but really it is no different from the problem that all models face" (p. 183). In any event, the fact that contextualists are so vocal and convincing in their rejection of associationism, and associationists are so silent or obscure in their treatment of context, suggests that context is nonassociative in nature. To show otherwise, PDP researchers must (a) be much more articulate in explaining the sort of incompatibility that contextualists such as Medin, Jenkins, and others (Bransford, Nitsch, and Franks, 1977) have perceived between context and patchwork spread of activation in discrete associative nets; and (b) show how contextual knowledge can emerge out of the sum of the connection weights among the units in an associative network. The same arguments apply to the use of other essentially nonassociative concepts such as *schema* (Bartlett, 1932; Iran-Nejad, 1980), *dynamic* (Iran-Nejad and Chissom, 1992), or *distributed* in connection with PDP connectionism (Iran-Nejad and Ortony, 1984).

Historical Foundations of PDP

In addition to associationism, the PDP perspective takes stock in formal computation. In fact, in the PDP literature, the statement that "knowledge is stored in . . ." may be completed with either (a) "associative connections" or (b) "numerical connection weights," interchangeably. In fact, traditional associationism and formal computation had already been brought together in the work of McCulloch and Pitts (1943) who devised a propositional calculus for representing the activity of what they called *nervous nets*. The most basic component of the nervous net was the reflex arc, an anatomical concept that represented a cyclical path of nerve associations which, "starting in some part of the body, passed by one way to the central nervous system, whence it was reflected over another to that same structure in which it arose and there inhibited or reversed the process that had given rise to it" (McCulloch, 1965, p. 266).

The work on artificial nerve nets by McCulloch and Pitts (1943) and Hebb (1949) gave rise to more than two decades of research on artificial neural networks until it was literally abandoned as a result of a review by Minsky and Pappert (1969). Interestingly, Minsky himself had already successfully built in 1951, with Dean Edmonds, the first artificial "nervous system" (Bernstein, 1981). In spite of this notable success, the disillusioned

Minsky convinced many researchers that it was more productive to turn away from subsymbolic neural networks and toward the symbolic computationalism of the information processing approach. Renewed interest in nerve-net research began a decade or so later with the publication of the book by Hinton and Anderson (1981), which launched the PDP research out of disillusionment with artificial symbol processing machines, and continued to grow on the same foundation that Minsky and Papert identified as inherently problematic.

DLR Research and Associationism

Researchers such as Lashley (1929, 1950, 1951) and John (1967, 1972), who have played the pioneering role in laying the foundation for DLR research, have argued strenuously against associationism, have viewed evidence suggestive of DLR as evidence against associationism, and have proposed DLR as an alternative to associationism. By contrast, PDP connectionists turned to associationism mainly because they were experiencing disillusionment about their own work in the artificial intelligence and information processing psychology of the 1970s. It is not surprising, therefore, that little attention has been given by connectionists to the foundational DLR research, which has had a long, steady, and respectable empirical and theoretical history. Instead, as just noted, connectionists have based their models on the formerly abandoned work in neural-net associationism by analogy to the network-like physical appearance of the brain.

Practically every researcher who has contributed to the notion of DLR, in the sense used by Lashley, has also found associationism unacceptable, going back at least to Dewey's (1896) seminal critique of the reflex arc. Lashley himself (1929, 1951) found his data as evidence against associationism:

> The results are incompatible with theories of learning by changes in synaptic structure, or with any theories which assume that particular neural integrations are dependent upon definite anatomic paths specialized for them. Integration cannot be expressed in terms of connections between specific neurons. (Lashley, 1929, p. 176)

Like Dewey (1896), Lashley (1929, 1950, 1951), and Jenkins (1974), Bartlett (1932) also stressed the incompatibility between associationism and his approach, maintaining that "the past operates as an organized mass rather than as a [patchwork] group of elements each of which retains its specific character" (p. 197). Bartlett also prophetically predicted that "in various senses, therefore, associationism is likely to remain, though its outlook is foreign to the demands of modern psychological science" (p. 308).

Bartlett's (1932) main objection to associationism was that "it tells something about the characteristics of associated details, when they are associated, but it explains nothing whatever of the activity of the conditions by which they are brought together" (p. 308). By adopting a nonassociative approach to the exploration of the underlying conditions, Bartlett also suggested that the associative approach is unsuitable for such exploration. However, the study of the activity of the conditions by which associative details, or subsymbolic (microstructural) units as they are now called, bind together into (symbolic) associative structures is exactly what PDP connectionism purports to do.

Summary

PDP connectionism and DLR foundational research represent two separate and fundamentally different worlds. DLR concepts, including the term *distributed* itself, behave like oil in water in the context of PDP connectionism; PDP modelers have paid little attention to DLR foundational research; the historical roots of PDP connectionism had already been rejected by DLR pioneers; and many of the leaders who contributed to the notion of DLR have found it necessary to reject associationism. Ignoring the inherent incompatibility between these two perspectives (e.g., by mixing concepts from the two) is likely to be detrimental to understanding the nature of human learning and remembering.

PDP and the Cognitive Revolution: Some Nonissues

Thus, PDP did not result from a systematic examination of the DLR foundational research. Neither did it develop as a result of the normal growth of the research on natural or artificial neural nets. Rather, it emerged as a reaction to the problems experienced by the artificial intelligence (AI) and information processing psychology of the 1970s; and it represents a return to a neural-net approach that was once abandoned by the very same leaders who had made the most compelling contributions to it. However, while all of this seems highly suggestive to us, none of it can by itself uphold the conclusion that PDP connectionism is untenable and unsuitable for future DLR research, as we happen to think it is. Such a conclusion must stand on the foundation of a more in-depth analysis of the nature of the PDP itself and its relation to DLR.

To those who closely witnessed the sudden shift of interest in mid 1980s from conventional to PDP cognitive science, the development may have seemed like an unusual turn of events, given what philosophers of science have been saying about the growth of scientific thinking. What is unusual is that some of the same researchers and many of the same concepts that played a direct role in popularizing the cognitive revolution in mid 1970s also played the leading role in promoting the PDP revolution in mid 1980s. However, the issue of whether the PDP approach is fundamentally different from traditional cognitive models is far from settled (see Oden, 1988; Pinker and Mehler, 1988). One source of confusion is the reluctance of the leading PDP connectionists to recognize the key areas of contrast between their own old (or conventional) and new (or PDP) approaches. Instead, they have often focused in their comparisons on notions such as the following that can readily be shown to be nonissues with regard to the phenomena under consideration.

Rules Versus Regularities — A True Dichotomy?

In 1975, Rumelhart introduced a set of syntactic and semantic story grammar rules for the structural analysis of narratives. These rules, devised by analogy to Chomsky's (1965) transformational grammar for sentences, were "to account for a substantial range of phenomena related to the higher order structures found in stories" (p. 234). Rumelhart's paper started an influential line of research that lasted close to a decade until the publication of two devastating reviews (Black and Wilensky, 1979; Wilensky, 1983) literally ensured that "the mistake in the analogy to sentence grammar should by now be obvious" (Wilensky, 1983, p. 582).

Later in 1986, McClelland, Rumelhart, and Hinton alluded to the story grammar experience stating that their PDP approach is completely different from "the 'explicit rule formulation' tradition, as represented by the work of Winston (1975), the suggestions of Chomsky, and the ACT model of J.R. Anderson (1983)" (p. 32):

> First, we do not assume that the goal of learning is the formulation of explicit rules. Rather, we assume it is the acquisition of connection strengths which allow a network of simple units to act *as though* it knew the rules. Second we do not attribute powerful computational capabilities to the learning mechanism. Rather, we assume very simple connection strength modulation mechanisms which adjust the strength of connections between units based on information locally available at the connection. (McClelland, Rumelhart, and Hinton, 1986, p. 32)

Rumelhart and McClelland (1986) have offered an influential demonstration of how PDP networks can manifest rule-like behavior without containing any explicit rules. They used the domain of the past tense of English verbs to show that rule-like behavior, similar to those used by children, can emerge out of nonrule-like underlying *regularities* without any explicit deep-structure rules or any other type of symbolic representations of verbs, roots, or suffixes. The basic network, called a *pattern associator*, consists of two layers of units: an input layer whose elements are connected to every element of an output layer by means of excitatory or inhibitory connections. The pattern associator learns by modulating the input–output connection weights using Rosenblatt's (1962) *perceptron convergence rule*. (Note that the pattern associator does not act *as though* it knows a rule; it blindly follows one.) On a given trial, the pattern associator takes an input pattern representing a verb stem; it calculates an output pattern representing its own version of a past tense; it compares this version to the correct version provided by an external "teacher"; and it adjusts its internal connection weights before the next trial. The pattern associator knows the association between the verb root and its correct past tense form when the actual output it calculates matches the correct feedback given by the teacher.

The rule-versus-regularity dichotomy crumbles in one's hands upon the most cursory examination. First, the PDP approach has itself produced a new generation of rules — and most aptly called such: the Hebb rule, the delta rule, the perceptron convergence rule, the back propagation rule, and so on. It is difficult to see the fundamental difference (at least along the rule-versus-regularity dimension) claimed by PDP connectionists between these rules and transformational rules (Chomsky, 1965), production rules (Anderson, 1983), story grammar rules (Rumelhart, 1975), or other types of rules postulated in conventional cognitive models. Secondly, the related explicit-versus-implicit distinction made between PDP connectionism and conventional approaches is equally fragile. For instance, following Polanyi (1958), traditional cognitive scientists routinely distinguish between explicit and tacit rules. In short, if there is any aspect that sets apart PDP rules from conventional rules, it has not yet been made clear.

The validity of any kind of rule depends first on how well it works. Rumelhart and McClelland's work on learning the past tense of English verbs has been carefully evaluated by linguists and other researchers (see Pinker and Prince, 1988). The general conclusion is that it is almost certain that language learning does not work in this way (Prince and Pinker, 1988). Nevertheless, as Prince and Pinker put it, PDP connectionists continue to "suggest, and many are quick to agree, that this shows the viability of associationist theories of language acquisition, despite their virtual abandonment by linguists 25 years ago" (1988, p. 195).

PDP and Behaviorism

There is even some question as to whether PDP connectionism can handle the mentalistic phenomena that differentiate conventional cognitive science and the strict behaviorism of the 1950s. This suggests that the standard contrasts made between the conventional cognitive science and PDP connectionism (those that we are identifying as nonissues) may themselves have been misguided to begin with. The inability of the PDP connectionism to deal with key mentalistic concepts such as attention or awareness along with its constant struggle to disassociate itself from stimulus–association–response (i.e., PDP's input–connection–output structure of the pattern associator or other connection machines) indicates that PDP connectionism is more similar to behaviorism than to cognitive psychology. Rumelhart and McClelland (1986) tried to dissociate PDP connectionism from behaviorism by claiming that PDP models allow hidden units and internal representations. This argument is not very convincing because behaviorists also postulated internal or hidden stimulus–response connections (short of postulating mentalistic concepts), which is exactly what PDP internal representations are: patterns of connection weights in the black box.

The correspondence between the way a pattern associator learns and the way a rat is said to learn in a Skinner box is striking. In both cases, there is a stimulus, a response, a reinforcement, and modification of the likelihood of the correct response. The difference is that in the pattern associator the stimulus and the response are explicitly identified as collections of subsymbolic (or uninterpretable) elements, as opposed to the stimulus–response elements that are implicitly uninterpretable. In both cases, the resulting output (the pattern or the response) is interpretable. Therefore, at least in this case, the subsymbolic–symbolic distinction reduces to semantics and carries no real substance. It is also not difficult to trace the historical roots of the PDP approach to radical behaviorism (see Reece, 1987). The perceptron model, on which the pattern associator is based, emerged out of strict behaviorism and disappeared with it. It is common among PDP authors to argue that the Minsky and Pappert (1969) book is responsible for the misfortune that befell perceptron research. Although this may, strictly speaking, be true, the explanation is too straightforward to capture the depth of the issues involved. It does not explain what prevented researchers from using the sort of handcrafted solutions found nowadays in the PDP literature, especially since many of these solutions were suggested by Minsky and Pappert themselves. A more compelling explanation might be that Minsky and Pappert's (1969) book had the same effect on perceptron-type "behaviorism" as Chomsky's (1959) review of Skinner's *Verbal Behavior* had on behaviorism in general, and for essentially the same reasons. Those reasons are still as alive and well today (e.g., in the work of those who find PDP connectionism unsuitable for language learning) as they were when Jenkins (1974), Lashley (1951), Bartlett (1932), Dewey (1896), and the Gestalt school argued against associationism.

Symbolic versus Subsymbolic

The PDP theorists also differentiate their models from conventional cognitive models along a symbolic–subsymbolic dimension. This is not exactly a nonissue (see Iran-Nejad, 1980, 1987); but it amounts to one in PDP connectionism. As the distinction goes, the conventional approach is the symbolic paradigm. Concepts and subconcepts serve as interpretable representations for external objects and their parts; and rules are interpretable tools of symbol manipulation. PDP connectionism is said to be a subsymbolic cognitive paradigm. The mind is said to work, metaphorically speaking, like the brain. Meaningful mental units and associations are replaced with nonsense units and connections; and mindful rules are replaced with mindless ones. These entities are said to serve, in large numbers, as subsymbolic representations of stored symbolic knowledge.

The symbolic–subsymbolic distinction is a major source of confusion with regard to how PDP and conventional paradigms are different. One problem is that subsymbolic entities in PDP connectionism are defined as interpretable mental phenomena: a unit "represents a hypothesis of some sort (e.g., that a certain semantic feature, visual feature, or acoustic feature is present in the input)" [Rumelhart, Smolensky, McClelland, and Hinton, 1986, p. 8]. Note that subsymbolic entities are not identified as features. Rather, they are *hypotheses* about features. Features could be uninterpretable; but it is difficult to imagine uninterpretable hypotheses. The problem sinks even deeper into the realm of incomprehensibility. The fact that terms such as *feature* have been used all along in the conventional symbolic paradigm — much in the same sense as they are used in the PDP subsymbolic paradigm (i.e., as elemental components of more complex structures) — is one of the major reasons for the controversy that subsymbolic connectionism might not be any different at all from the conventional symbolic perspective.

Computer-Inspired versus Brain-Inspired

Another aspect of the conventional cognitive science that PDP modelers often compare to their new connectionism is that the conventional approach is computer-inspired as opposed to brain-inspired. Two interrelated aspects are generally mentioned in this regard: speed and parallel processing. The computer is fast and sequential, while the brain is slow and parallel. Although these might be important considerations in AI, they are nonissues in cognitive science. First, the brain is capable of sequential symbol processing (SSP), like or unlike computers. But since PDP models provide no clues whatsoever as to how the brain engages in SSP, there is no basis for finding them to be more brain-inspired. Secondly, the fact that conventional computers are incapable of parallel processing — and the brain is — does not mean that the conventional cognitive science goes out of the window. Both subsymbolic and symbolic aspects of cognition must be explained, preferably simultaneously in terms of a single paradigm (see Bereiter, 1991). It is a basic tenet of PDP connectionism that what makes the brain fast and efficient is *subsymbolic* parallel processing. However, one might assume, on the contrary, that the brain — especially that of human beings (as compared to animals) — is capable of *symbolic* parallel processing (SPP), and that is what makes the brain fast and efficient.

Many cognitive scientists now agree that the brain and the computer do things differently. Many also agree that conventional cognitive science probably went too far in taking the computer metaphor literally. However, the new PDP connectionism has yet to show exactly where cognitive science went wrong in making use of the computer metaphor as a model of human cognition.

Similarly, few researchers would dispute the fact that "sooner or later, theories of cognition will have to deal with the problem of the relationship between the neuronal network and the conceptual network" (Iran-Nejad, 1980, pp. 10–11). However, PDP connectionists have engaged in what amounts to the circumvention of the mind–brain problem. Specifically, PDP connectionists used the term *brain-inspired*, instead of brain-based, to avoid the problem of directly addressing the mind–brain problem. The term *brain-inspired* licenses PDP connectionists' metaphoric adoption of neurophysiological concepts such as activation, inhibition, threshold, or summation. It also assures the freedom from having to deal with the real nature of "the relationship between the neuronal and the conceptual networks [that] must be theoretically clarified before neurological concepts can be used in the psychological domain" (Iran-Nejad, 1980, p. 12).

Given the associative assumptions and the "overly complicated picture" cognitive models such as those postulated by Rumelhart (1975, 1977, 1978) portray, the issue of a bridge between mental phenomena and the nervous system seems too remote to even consider (Iran-Nejad, 1980). The willingness to use the brain as a metaphor for thinking about the microstructure of cognition has enabled the PDP modelers to use some brain concepts (e.g., neural networks, activation, inhibition), while at the same time keeping their distance from the brain itself. This is, of course, a classic case of having one's cake and eating it as well. Brain concepts are as alien in the PDP context as the term *distributed* is in the context of associationism. Therefore, PDP connectionists have little advantage over conventional models as far as the study of the mind–brain relationship is concerned.

Summary

It is not exactly clear how cognitive rules emerge out of nonrule-like computation. However, PDP rules such as the delta rule are still rules. Parallel distributed processing connectionism is more comparable to behaviorism than to cognitive science. The symbolic–subsymbolic dimension is too ill-defined to differentiate the PDP and conventional approaches. Connectionism may be called brain-inspired, but it contributes very little to our knowledge about the mind–brain relationship. In other words, many of the areas that PDP connectionists have selected as the basis for comparing their models with conventional approaches only have the appearance of being important issues. Once one tries to step beyond rhetoric, the seemingly important issues quickly change into nonissues.

Distributed Learning and Remembering: Some Key Issues

The Computer Metaphor and the Autonomy of the Product

The computer software metaphor was explicitly introduced by Neisser (1967) who used the program analogy to argue that long-term mental schemas can literally exist in the head. He argued that the study of mental programs is what psychology is all about. This was the beginning of the information processing cognitive psychology that viewed the mind as operating, like a computer, on a vast number of mental files. Discovering the organization of these files in the long-term storehouse, as well as the processes by which they were stored and retrieved, became the central goals of the cognitive science that reached the peak of its popularity in mid 1970s. Here is where, we believe, the cognitive revolution went wrong: it became exclusively the science of the processing, storage, and organization of static mental software.

The clean separation of the mental software from the brain wetware is a good example of how PDP and conventional models are identical in their fundamental assumptions, in spite of their superficial differences. Recall that PDP units are not brain units such as neurons; they are mental entities. In other words, the subsymbolic level is a purely mental-software level dealing exclusively with the analysis of the microstructure of cognition. No consideration whatsoever is given to the functional properties of the brain. It is assumed that the subsymbolic software of the mind is formalizable as an autonomous system in its own right (Smolensky, 1988). In short, like traditional cognitive models, which were proposed by some of the PDP leaders in mid 1970s, the PDP models born in mid 1980s are concerned with the study of the mental software. According to McClelland, Rumelhart, and Hinton (1986):

> They [PDP models] hold out the hope of offering computationally sufficient and psychologically accurate mechanistic accounts of the phenomena of human cognition which have eluded successful explication in conventional computational formalisms; and they have radically altered the way we think about the time course of processing, the nature of representation, and the mechanisms of learning. (p. 11)

This quotation claims that PDP models are paradigmatically different from traditional cognitive models. However, it also shows that the two approaches share the fundamental assumption of the autonomy of the mental software. Like pure mathematics, pure cognition can also exist.

Is the analysis of the features of mental products (or software) likely to shed light on our understanding of how people think and solve problems? An analogy might be helpful here. Imagine a group of researchers preoccupied with the analysis of the curves, lines, angles (and so on) of the photographs a camera takes. The goal is to understand the camera's method of storing and reproducing the photos. Think for a moment about the analysis of such features, their optimal organization in an accumulated corpus of photos, and the manner by which they interconnect. How likely is such an exploration to shed any light at all on the picture-taking capacity of the camera itself? The photograph of an elephant is different from that of an octopus. Do these differences have anything to do with the picture-taking capacity of the camera? This analogy was once used to illustrate the inherently problematic nature of conventional cognitive science of the 1970s (Iran-Nejad, 1980). It applies equally well to PDP connectionism despite the brain-inspired rhetoric.

The Schema: A Long-Term Structure or a Transient Functional Pattern

A key concept in conventional cognitive science of the 1970s was the notion of schema. Rumelhart (1980) defined the schema as a long-term memory monolith — as the most elemental building block of cognition. Long-term memory schemas made it difficult to explain how people reorganize their thinking in order to recontextualize their understandings from one external setting to another. Iran-Nejad (1980) used the comprehension of the surprise-ending story to illustrate the problem and concluded that, in order to solve this problem, schemas had to be viewed as transient ongoing patterns (see Iran-Nejad and Winsler, 2000, this issue).

Iran-Nejad (1980) used a light bulb constellation analogy to illustrate how the brain might create and uphold transient schemas without storing permanent representations of them. Two basic assumptions were made about the brain and its functioning. The first assumption was that the brain is populated with a large number of living microsystems. In terms of its self-reflective discriminability, each microsystem was assumed to be analogous to a color-coded light bulb. The second assumption was that knowledge was none other than the "live" self-awareness of the activity of ongoing constellations of the brain microsystems (or neurons). Not only had the light bulb analogy spawned a new conception of knowledge and knowledge schemas (as ongoing patterns of self-awareness); it had also quite naturally introduced the notion of biofunctional distributed learning and remembering (BDLR), which we discuss later in this article.

The transient-schema hypothesis gave rise to more questions than answers, the most urgent of which had to do with the nature of long-term remembering. Nevertheless, it made it possible to draw a rough picture of how the brain is capable of spontaneous recontextualizations of one's thinking from one moment to another in terms of what Bartlett (1932) called the (ever-evolving) *schema-of-the-moment.* The surprise-ending story illustrated the most dramatic instance of the spontaneous recontextualization power of the brain. It showed how people used the elements of one (pre-surprise) schema-of-the-moment to create a different (post-surprise) schema-of-the-moment. Over the years, some of the initial questions surrounding the transient-schema hypothesis have been addressed, including those having to do with how the brain engages in sequential and parallel symbolic processing (Diener and Iran-Nejad, 1986; Iran-Nejad, 1989a, 1989b, 1989c; Iran-Nejad, Clore, and Vondruska, 1984; Iran-Nejad and Ortony, 1984).

Levels of Analysis

The concept of levels of analysis has become an area of concern for PDP models ever since Broadbent (1985) posed the problem in his commentary on McClelland and Rumelhart's (1985a) distributed model. Broadbent suggested that the proper level to consider the hypothesis that memory is distributed is the physiological level (see Iran-Nejad, 1980) rather than the psychological level. He therefore, questioned the validity of the psychological evidence McClelland and Rumelhart used to compare their model with traditional cognitive science models they examined.

Broadbent reasoned that at a purely psychological level distributed and localizationist approaches are indistinguishable from each other. To be sure, a theory postulating localized (unitary) psychological phenomena may be either distributed or localized at the level of the physical brain; and an approach maintaining that memory is distributed at the physiological level is compatible with one postulating localized entities at the psychological level. Broadbent concluded that the evidence to which McClelland and Rumelhart "appeal is from a different level of explanation, and therefore irrelevant to the undoubted merits of the distributed approach" (p. 192).

In their rejoinder, McClelland and Rumelhart (1985b) pointed out that "Broadbent's analysis went astray" because it was based on an incomplete consideration of Marr's (1982) notion of levels. Broadbent's (1985) discussion referred to Marr's implementational (or physiological) and computational (or psychological) worlds. It did not mention Marr's algorithmic (still psychological) world. McClelland and Rumelhart reasoned that "indeed, it would appear that this [i.e., the algorithmic level of analysis] is the level to which psychological data speaks most strongly" (Rumelhart and McClelland, 1986, p. 123). And this was exactly the point Broadbent was rejecting: whereas the evidence in support of the notion of distributed remembering would only be meaningful at the physiological level of analysis, the psychological data such as those presented by McClelland and Rumelhart applied only to psychological levels and, therefore, would have no bearing on the DLR hypothesis.

More specifically, McClelland and Rumelhart (1985b) failed to consider two aspects of Broadbent's (1985) discussion. First, Broadbent's main point was that the term distributed was applicable to the level of the physical brain (i.e., to the level of the brain hardware, so to speak) and not to any purely psychological (or software) level — algorithmic or computational. Secondly, McClelland and Rumelhart did not consider the possibility that in using only two levels, Broadbent was being guided by the concept of natural levels without getting involved with the proliferation of artificial levels handcrafted arbitrarily by investigators. After all, why only two cognitive levels and not more.

By instantiating their model more firmly at a purely psychological level, McClelland and Rumelhart (1985b) did not solve the problems raised by Broadbent. Instead, they raised another issue more problematic for the PDP approach: whether or not the term *distributed* can be applied at all to the psychological level. The issue became substantially more complicated in a recent article by Smolensky (1988) who claimed that PDP units and connections have no spatial location, as opposed to neurons and synapses that are located in three-dimensional space. In any event, at least one thing is certain in considering the PDP discussions of the notion of levels: Lashley's (1929, 1950, 1951) evidence, which was gathered based on the assumption of distributed learning and remembering in the three-dimensional brain, has little, if any, bearing on the meaning of the term *distributed* used in the expression *parallel distributed processing.*

Intralevel versus Interlevel Approaches

In the expression *distributed learning and remembering,* the term *distributed* makes more concrete sense at the physical level of the three-dimensional brain as was originally conceived by Lashley (1929). However, the physical brain is not the proper domain for psychological research. Learning and remembering, on the other hand, are more directly psychological than physiological problems. One way to take the dilemma out of the notions of PDP or DLR is to use the brain as a metaphor for the development of a conceptual framework for psychological research: to "replace the 'computer metaphor' as a model of mind with the 'brain metaphor' as model of mind" (Rumelhart, Hinton, and McClelland, 1986, p. 75). This is the route PDP connectionists took in search of a home for the concept of distributed representations at a purely psychological level. They postulated a hypothetical subsymbolic level and selected that as the proper domain for solving psychological problems. In order to be able to dispose of the cake of the brain and keep it as well, PDP connectionists introduced the subsymbolic psychological level using the brain as a metaphor for thinking about the microstructure of cognition:

> Though the appeal of the PDP models is definitely enhanced by their physiological plausibility and neural inspiration, these are not the primary bases for their appeal to us. We are, after all, cognitive scientists, and the PDP models appeal to us for psychological and computational reasons. (McClelland, Rumelhart, and Hinton, 1986, p. 11)

However, if the arguments reviewed so far are correct, the cake may not exist at all. There is nothing to eat or keep. PDP connectionism faces several problems. It reduces cognitive science to (parallel) behaviorism; the term distributed is not applicable to the psychological level, literally or metaphorically; and the assumption that there exists a microstructural level at which cognitive phenomena may be analyzed and mathematized autonomously is untenable.

If the term *distributed* is interpretable only in reference to the physical and physiological brain, as Broadbent (1985) argues, then exactly how is it relevant from a psychological point of view. Clearly, psychologists cannot choose neuroanatomy and physiology as the proper domain for their research. In what sense then is the evidence supplied by Lashley (1929), who systematically removed sections of the brains of laboratory animals and examined their learning and retention capabilities, applicable to the psychological level?

There is a solution to this dilemma that is compatible both with Broadbent's (1985) view and Lashley's (1929) evidence: to abandon the assumption of the *autonomy of the product* as a separate level of analysis and view psychology as an *interlevel* discipline (Iran-Nejad, Clore, and Vondruska, 1984). More specifically, confining one's hypotheses and theories to the phenomena at the purely psychological level or the purely physical level may be referred to as an *intralevel* approach. Both conventional cognitive models and the PDP models are intralevel approaches. They share the assumption of the autonomy of the product and focus directly on the analysis of the features of the product. Lashley, on the other hand, was conducting interlevel psychological research. He was testing hypotheses about learning and remembering in terms of the functional properties of the physical brain.

Because intralevel approaches explain more complex phenomena in terms of simpler entities at the same level, they necessarily encounter the problem of reduction (Wimsatt, 1976). PDP models, for instance, explain complex mental schemas in terms of associative connections among large sets of mental microfeatures. The units, the connections, and the whole schema are all in one and the same (cognitive) level. Therefore, neither of the approaches can solve the problem often stated in terms of the Gestalt dictum that the whole is more than the sum of its parts.

The interlevel approach explains cognitive phenomena such as learning and remembering in terms of the relevant functional properties of the brain, as did Lashley (1929, 1950, 1951). Functional properties of the brain and its components, then, become the proper domain of psychological research and the source of psychological hypotheses. The assumption of the autonomy of the product, held by Marr (1982), PDP connectionists, and conventional cognitive scientists would have to be abandoned altogether. This does not mean, however, "that in order to learn about cognition, one would have to open the head [as did Lashley] and directly examine the neuronal organization" (Iran-Nejad, 1980, p. 29). Like Lashley (1929), Bartlett (1932) took a functional approach, but without working directly with the physical brain. What distinguishes psychology from neurophysiology is the type of questions addressed rather than simply differences in research methodology. To use the camera analogy, we must try to understand how the camera works as an integrated system, instead of directly analyzing and formalizing the features of the pictures the camera takes. We can then continue to test our hypotheses about the functional properties of the camera and its component parts by manipulating variables related to the different types of pictures the camera takes.

This leads us back to the distinction between schemas as monolithic long-term memory structures capable of being stored, like computer programs, autonomously of the information processing system that uses them (Rumelhart, 1978) and the view that they are live patterns of awareness capable of existence only so long as they are being created and upheld by the ongoing activity of the brain, functioning as a whole. Iran-Nejad (1980) traced this structural–functional distinction to the work of Bartlett (1932) who argued that the functional level, and not the structural level, is the proper domain of experimental psychology.

The hypotheses that the brain is populated with a large number of microsystems capable of generating self-awareness and that knowledge is live self-awareness created by the ongoing activity of brain microsystems are interlevel hypotheses. They are about psychologically-relevant biofunctional properties of the nervous system. Lashley's (1929) evidence is directly relevant to these hypotheses and to other aspects of the biofunctional model because such a model can address the issue of how the functional properties of distributed constellations of brain microsystems can create and uphold ongoing mental schemas.

Summary

Like conventional approaches, PDP connectionists maintain the assumption of the autonomy of mental products. They analyze, autonomously of the functional properties of the nervous system, complex mental and behavioral structures into mental microunits and connection weights. As a result, they face the insoluble problem of reduction which is characteristic of all intralevel approaches. An alternative way of thinking about distributed learning and remembering was proposed based on the assumptions that mental structures are nothing other than live self-awareness, they can exist only while they are being created and upheld by the brain, they are directly unanalyzable, and they can be explained in terms of the functional properties of the brain and its components (Iran-Nejad, 1980). It is only by taking an interlevel approach of this kind to the study of the mind that psychologists can avoid the problem of reduction (Iran-Nejad, Clore, and Vondruska, 1994).

Biofunctional Distributed Learning and Remembering

As already discussed, research on biofunctional distributed learning and remembering (BDLR) began with an interest in the nature of mental schemas. The DLR aspect, although an inherent aspect of the model, has not been directly addressed in any detail in the past. The immediate goal of the biofunctional approach was to examine mentalistic concepts such as schemas, learning, awareness, attention, and remembering in terms of the functional properties of the brain (Iran-Nejad and Ortony, 1984), without reducing these concepts to nonmentalistic brain concepts such as activation, inhibition, and connection weights (Iran-Nejad, 1980).

A detailed analysis of the biofunctional schema theory of learning and its relationship with modern schema theories has been given in Iran-Nejad and Winsler (2000, this issue). The relationship of the model with PDP schema theory (Rumelhart, Smolensky, McClelland, and Hinton, 1986) has also been already briefly examined, both theoretically (Iran-Nejad, 1989a, 1989b) and empirically (Iran-Nejad, 1989c). In this last section, we will focus on the DLR aspect with an emphasis on its relationship to Lashley's (1929, 1950, 1951) research.

The basic biofunctional assumptions directly bearing on DLR will be presented first. It is useful to bear in mind that learning and remembering may be viewed as distributed in the sense that they involve brain activity occurring simultaneously in constellations of neurons located physically at various parts of the brain. There is, however, another related sense in which learning and remembering may be viewed as distributed: the factors that contribute to learning are themselves distributed in the sense that they come from multiple sources (Bereiter, 1985; Iran-Nejad, McKeachie, and Berliner, 1990).

Brain Components as Subsystems and Microsystems

There is a tendency among some researchers to think of the brain as an anatomical mass, a mushroom of areas, centers, or regions packed against one another and connected by a network of neural pathways. This anatomical mass is thought to house the mind in a hierarchical fashion, with simple sensory materials stored statically at lower levels and complex mental schemas at higher regions. We believe that this view of the brain is almost certainly wrong.

It is more likely that the nervous system comprises an indefinite, but not inordinately large, number of interrelated subsystems, of which sensory subsystems are but a subset. Similarly, the smallest unit of the nervous system is probably not a nerve fragment, a chemical particle, or a synapse, but a functionally autonomous microsystem, equivalent perhaps to what is traditionally called the *neuron*. Every subsystem would, then, consist of a vast number of microsystems. Brain subsystems and microsystems do not store static mental entities of any kind. They create and uphold them in the form of live, as opposed to pre-recorded, self-awareness from one moment to another.

Heterogeneous and Homogeneous Specialization

The microsystems in different subsystems are specialized, phylogenetically or ontogenetically, to create qualitatively different subjective experiences. An obvious example of this is that the microsystems responsible for audition create self-awareness of sound, which is a qualitatively different experience from the self-awareness created by the microsystems involved in vision. Thus, with regards to the subjective experiences they create, the microsystems in different subsystems are heterogeneously specialized. This sort of heterogeneous specialization is assumed to play the primary role in causing diversity in subjective experiences. This diversity is evolution's solution to the problem of complexity.

The microsystems within each subsystem, on the other hand, are homogeneously specialized. This simply means, for instance, that all the microsystems within the auditory subsystem are similarly specialized to create auditory experiences, although, of course, there is enough variety among them to deal with all the different aspects of auditory discrimination including those involved in the perception of speech (e.g., pitch, rhythm, intonation).

Group Distribution and Constellation Distribution

The term *distributed* in *BDLR* has two related meanings. Sometimes mental entities are created and upheld by a group of localized microsystems (group distribution), much in the same way that the traffic arrow on the freeway is created by a group of adjacent light bulbs. Group distribution is primarily a within-subsystem phenomenon and is responsible for the creation of such basic experiences as pure auditory and visual imagery. The display on your calculator forms numbers, your computer screen displays letters and words, and your television screen shows pictures by taking advantage of the principle of group distribution, although they all do so in a purely mechanical fashion.

The "mental displays" formed by group distribution may be described as a one-to-many type of relationship — where *one* refers to the mental product, and *many* refers to the brain microsystems. To use the traffic arrow analogy, the arrow display (one) and the blinking light bulbs (many) creating and upholding it constitute an example of such a relationship. It is important to reiterate that this is an interlevel relationship with the term *one* referring to the mental display and the term *many* referring to the group of brain microsystems creating the display. At least in the biofunctional model, this has to be the case necessarily because the model abandons the assumption of the autonomy of the product. No static trace of the arrow remains, in any form whatsoever, when the light bulbs go off. It is rather obvious here that it makes no more sense to try to account for the construction of mental displays by binding together a group of smaller feature-like entities any more than it makes sense to build a traffic arrow by directly patching together a number of light spots (Iran-Nejad, Clore, and Vondruska, 1984).

The interlevel one-to-many relationship (Iran-Nejad, 1980; Iran-Nejad, Clore, and Vondruska, 1984) must, therefore, be distinguished from the intralevel one-to-many relationship more recently postulated by PDP theorists. For instance, Rumelhart, Hinton, and McClelland (1986) pointed out that the one-to-many distributed representations they postulated, in which units represented small, feature-like entities, "should be contrasted to a *one-unit-one-concept* representational system in which single units represent entire concepts or other large meaningful entities" (p. 47). The PDP one-to-many distributed representations are intralevel because the terms one and many both refer to mental entities.

The biofunctional (interlevel) one-to-many relationship must also be distinguished from the (intralevel) one-to-many connection patterns evident in the anatomical layout of the nervous system (Hinton and Anderson, 1981). For instance, referring to the anatomical structure of the nervous system, Broadbent (1985) stated that "the *one-to-many* and the *many-to-one* connections of the sensory pathways resemble much more closely the [distributed] diagrams constructed by Winograd and Cowan than they do the specific paths needed for the alternative form of coding" (p. 190). Here again Broadbent is referring to an intralevel one-to-many relationship because both ends of the one-to-many relationship are located at the same anatomic level, where one (complex unit) is anatomically connected to many (simple) units. PDP theorists adopt this apparent, one-to-many, neuroanatomic architecture as a metaphor for postulating group distributions at the level they call the microstructure of cognition.

Group distribution is the only sense in which PDP connectionists use the term *distributed*. From the biofunctional perspective, within-subsystem group distribution can do little more than generating meaningless mental displays. Meaningful mental phenomena such as concepts and schemas require the participation of many subsystems. And this is where constellation distribution enters the stage: meaningful mental entities are created and upheld by widely-distributed constellations of brain microsystems, whose members are physically located in many different subsystems. The brain creates the concept of *dog*, for instance, by means of a distributed constellation of microsystems whose members are distributed across the visual, auditory, affective, and many other subsystems (Iran-Nejad and Ortony, 1984). The biofunctional model takes advantage of the notion of constellation distribution to explain Lashley's (1929) findings.

Equipotentiality in the Nervous System

Lashley (1929) used the term equipotentiality to "designate the *apparent* capacity of any intact part of a functional area to carry out, with or without reduction in efficiency, the functions which are lost by destruction" of another area (p. 25, italics added). PDP connectionists and others have pointed out that Lashley's distributed theory went too far in its assumption of equipotentiality in brain functioning. This we believe is too quick an evaluation. In fact, we will try to show in this section that by hedging his definition of equipotentiality with the term *apparent*, Lashley probably did not go far enough.

According to Lashley (1929), "the contribution of the different parts of a specialized area . . . is qualitatively the same" (p. 176). This means that the parts of a subsystem are equipotential in their ability to create and uphold mental displays of the patterns for which they are specialized. The computer screen on which we are typing this paper is equipotential in the sense that any portion of it can display any specific letter or letter sequence as readily as any other portion; the different portions of your television screen are equipotential because any portion can display any picture that other portions can display; and so is the electric sign described by Lashley: "The activity of the visual cortex must resemble that of one of the electric signs in which a pattern of letters passes rapidly across a stationary group of lamps. The structural pattern is fixed, but the functional pattern plays over it without limitation to specific elements" (1929, pp. 158–159; see Iran-Nejad, 1980, 1987 for further discussion of the differences between structural and functional patterns).

It is important to bear in mind that Lashley's (1929) concept of equipotentiality should not be contrasted with specialization in the nervous system. Consider, for instance, within-subsystem specialization. Your color television screen reflects much more specialization than your black-and-white television screen. Yet, both screens are equally equipotential in the sense that they can display any specific picture anywhere on them. The analogy can be extended meaningfully to a color-blind visual system. What equipotentiality should be contrasted with, therefore, is the notion of fixed specific traces. If particular letters or letter sequences could only be produced on specific parts of a computer screen, that screen would be of little, if any, value. Likewise, if certain images could only be displayed on specific sections of a television screen, that television would be equally worthless.

One cannot ignore the possibility that the PDP patterns that store specific schemas in their connections are equally worthless. Earlier in this paper, we mentioned that Lashley's (1929, 1950, 1951) nonassociative approach is incompatible with the associative approach taken by PDP connectionists. We can now be more explicit about this claim. The idea is that Lashley's concept of equipotentiality and the notion of specific traces are incompatible, whereas associative structures are necessarily synonymous with specific traces:

> A pattern of activation only counts as the same as another if the same units are involved. The reason for this is that the knowledge built into the system for re-creating the patterns is built into the set of interconnections among the units . . . . For a pattern to access the right knowledge, it must arise on the appropriate units. In this sense the units play specific roles in the patterns. (McClelland and Rumelhart, 1986, p. 175)

A great deal has been made of the evidence collected by Hubel and Wiesel (1968, 1977) which purports to show that there are specific feature detectors in the primary visual cortex. There are simple cells that respond specifically to lines or edges of a particular orientation in a specific place in the visual field. Based on this and similar evidence, McClelland and Rumelhart (1981) postulated a hierarchical model of word perception. Feature detectors at the lowest level are connected to letter detectors that are, in turn, connected to word detectors. Iran-Nejad and Ortony (1984) used what was considered to be a simple but highly informative example of between-subsystem equipotentiality to argue against the notion of specific detectors in the nervous system:

> Blindfolded subjects are capable of correctly identifying letters (finger-written) on their skin. White, Saunders, Scadden, Bach-Y-Rita, and Collins (1970) used a visual substitution apparatus which converted optical images into tactile displays which blind or blindfolded subjects were able to "see with their skin." It was shown that "subjects are able to perceive certain simple displays . . . almost as soon as they have been intro-

> duced" (p. 23) and that with minimal amounts of training they are able "to identify familiar objects and to describe their arrangement in depth" (p. 25). (Iran-Nejad and Ortony, 1984, pp. 182–183)

Not only does the tactile system perform like the visual subsystem, but it also shows a good deal of within-subsystem equipotentiality in doing so. Finger-written letters can be identified almost anywhere on the skin, although the "high resolution" — with regards to the number of tactile microsystems available — finger tips can do so much more readily. In fact, neuropsychologists use the failure to identify numbers written on the finger tips as evidence of damage to the intact nervous system. How does the tactile subsystem manage to perform like the visual system without any prior experience? Unless we are willing to acknowledge the existence of innate feature detectors, letter detectors, number detectors, and object detectors everywhere under the skin and beyond, we are forced to abandon the notion of specific traces and accept the hypothesis of between-subsystem equipotentiality.

It is interesting to note that the degree of between-subsystem equipotentiality varies from subsystem to subsystem. The visual subsystem is more equipotential to the tactile subsystem than to the motor subsystem while the auditory subsystem seems to be more equipotential with the motor subsystem than with the tactile subsystem. It is much more difficult to dance to the patterns created in the visual subsystem than it is to those created in the auditory subsystem. It is also possible that there exist subsystems within the nervous system that are equipotential to many or all other subsystems. In any event, as Lashley (1929) argued and his evidence indicated, between-subsystem equipotentiality is a matter of degree.

Equipotentiality of Neural Connections

The subsystems and microsystems of the nervous system are interconnected by a vast network of neural pathways. It is often assumed by neural-network connectionists (e.g., Sperry, 1943) that the neural network consists of specific element-to-element connections, each of which is associated with a connection strength (or weight) that determines the degree of activation (or influence) a source element can exert on a target element.

Based on biofunctional schema theory, Iran-Nejad and Ortony (1984) argued that there is another alternative. It is the specialization of the brain microsystems and not the specificity of neuroanatomic connections that is the basis for communication in the brain. Iran-Nejad and Ortony (1984) were able to trace this nonconnectionist alternative to the work of Weiss (e.g., 1936). Consider the following example: if a portion of the skin from the belly of a salamander is removed and planted on its back, and if, after regeneration, this skin (which is now located on the back of the animal) is stimulated, the animal proceeds to scratch its belly, the original location of the skin. Sperry (1943, 1966) explained evidence such as this by postulating some sort of chemical affinity between the peripheral cells (belly skin cells in this case) and corresponding central cells (those commanding the scratching behavior) that enables the former to trace their way back and reestablish specific contact with the latter. However, there is no need for the reestablishment of a new nerve line. According to the nonconnectionist hypothesis, all that has to be reestablished is contact with the local nerves and, thereby, with the nervous system as a whole. What makes the animal think that the belly skin is being stimulated is the specialization of the belly-skin cells themselves, regardless of their location so long as belly skin cells are connected to the neural network as a whole.

Iran-Nejad and Ortony (1984) argued that the neural network is an all-purpose communication system serving all brain subsystems and microsystems (equipotentially). In other words, it is not a collection of specific element-to-element connections representing associative links of varying degrees of (synaptic) strength. Rather, it operates much in the same way that the blood circulation system does — as an all-purpose network serving many bodily subsystems. Consider, for example, the case of the endocrine glandular subsystem. Endocrine glands release their products into the extracellular fluid surrounding capillaries. These hormones then enter the blood circulation system which carries them, as an all-purpose system, everywhere the blood goes and not directly to their specific target organs.

For instance, ACTH is released in the anterior pituitary gland, located on the lower part of the brain. This hormone stimulates the cortical adrenal cells located above the kidneys. A direct duct could have been used by the system to carry ACTH from its source to its target organs but, "if a tube were to be available from every endocrine gland to its target organ, organisms would become monstrously complex" (Iran-Nejad and Ortony, 1984, p. 182). A much more elegant way to solve the problem is to use the all-purpose, all-serving, and all-spreading blood circulation system. The hypothesis, proposed in the context of the biofunctional system, that the neural network is an all-purpose communication network is consistent with Lashley's (1929) notion of equipotentiality of neural pathways stating "that the interruption of association or projection paths produces little disturbance of behavior, so long as cortical areas supplied by them remain in some functional connection with the rest of the nervous system" (p. 175).

Learning and Remembering as Multisource Phenomena

Association, as the most basic unit of learning, has been a long-lasting assumption in psychology. In its most sophisticated form yet, associationism is gaining widespread acceptance in the form of PDP connectionism, which defines knowledge as connection weights among specific (localized) groups of units and learning as changes in those connection weights. The biofunctional distributed learning and remembering model, on the other hand, assumes that learning is not the establishment of specific connections. Rather, even the simplest meaningful act of learning involves the participation of many brain subsystems and many factors. In other words, the creation of image-like within-subsystem mental displays (i.e., group distribution activity) cannot be considered meaningful learning. The meaning of a mental display in a given subsystem is the combination of all the displays that occur in other subsystems simultaneously. The more subsystems involved, the more complex the meaning. This is one sense in which learning and remembering may be viewed as distributed.

Another sense in which the term *distributed* may be examined is that the factors that contribute to learning are distributed. In other words, according to the BDLR approach, external information is not the only source of learning. Learning is a multisource phenomenon (Bereiter, 1985) and by far the majority of the sources that contribute to any meaningful act of learning or remembering are internal sources (Iran-Nejad, 1989a, 1989b).

Among the major sources contributing to BDLR are those that regulate the functioning of the system. According to the multisource hypothesis, the locus of learning lies in the interaction of the sources of control and learning processes. Traditional discussions of control generally consider two sources: external and internal. Behaviorists assume that learning occurs passively under the control of external stimuli. Learning of this sort can hardly go beyond reactive attention, brief surprise states, momentary orientation to events, local (or within-subsystem) combination, categorical knowledge creation, and piecemeal metacognition. It is clear, therefore, that by itself external control can result in little, if any, meaningful learning.

The passive learning of the type postulated by behaviorism is often contrasted with active learning in conventional cognitive science. Active learning is said to occur under the conscious (or effortful) control of the individual learner represented internally as the central executive. In this sense, active control implies that the nonexecutive components of the system are passive and cannot contribute to learning without executive monitoring, that is, unless they are directly acted upon by the executive component. This can be illustrated most clearly by Neisser's (1967) original analysis of (re)construction in the human information processing system in terms of his dinosaur reconstruction analogy.

According to Neisser, people construct new schemas, or reconstruct old schemas, by piecing together the stored fragments of their past schemas in much the same way that paleontologists reconstruct the replica of a dinosaur by piecing together bone fragments of an extinct dinosaur. Neisser's analogy suggests that there is one and only one source that regulates the process of knowledge construction (i.e., the executive control), in the same way that there is one and only one source that regulates the process of dinosaur reconstruction (i.e., the paleontologist). The other components that participate in information processing construction (i.e., Neisser's abstract schema and the concrete content that fills the slots of the abstract schema) are static and passive in much the same way that the bone fragments and the mortar used in dinosaur reconstruction are static and passive.

The BDLR model, on the other hand, implies that it is highly unlikely that DLR is under the exclusive influence of active (or executive) control. It can, on the other hand, occur readily under the exclusive influence of dynamic control and optimally in a multisource fashion: under extensive dynamic, proper active, and optimal external control. Dynamic control is responsible for what is generally known as incidental learning, learning without knowing, and nonstrategic learning.

There are a number of ways to contrast the two internal sources of control. Active internal control is effort-mediated while dynamic internal control is interest-mediated. This aspect can be most clearly illustrated in terms of inquiry processes of the system. Under active internal control, inquiry processes take the form of strategic questioning based on effortful thinking. The more effort we apply, the more questions we end up asking ourselves or others. Under dynamic control, inquiry processes take the form of curiosity which is often viewed as being synonymous with interest (Berlyne, 1974). When multiple sources are at work, the person can sustain an exploring or investigative attitude.

Secondly, active and dynamic control may be contrasted in terms of intralevel and interlevel functioning of the system. Active control is intralevel in that it is encapsulated within the realm of mind-mediated strategies and plans. It is the type of control that is implemented effortfully and strategically by the person. Dynamic control is also intralevel at the level of the brain; subsystems and microsystems of the brain directly regulate learning processes. Dynamic control depends on the biofunctional properties. For instance, the more equipotential the brain subsystems (e.g., the auditory and motor subsystems as opposed to the visual and the motor subsystems) the more readily they tend to work together (e.g., as it happens in dancing to music). Multisource self-regulation tends to be interlevel, when the three sources of self-regulation combine forces in the brain–mind cycle of reflection (see Iran-Nejad, 2000, this issue). The idea here is that the individual gains control over the components of his or her nervous system (i.e., its subsystems and microsystems) much in the same way as he or she has control, for instance, over his or her own limbs and fingers.

Under many circumstances, dynamic control follows active control so naturally that the contribution of the former is taken for granted. But it is not difficult to show that execution of active control is likely to be nearly impossible without the contribution of dynamic control. First an analogy. Imagine trying to steer a car with power-steering when the power is off. Normally, the individual driver exerts active control; but the power in the steering provides something analogous, roughly speaking, to dynamic control. When dynamic control is off, active control is difficult if possible at all. A situation very similar to power-steering is trying to move one's arm when the arm is asleep. Since the individual is awake and can readily control other parts of the body, the active control system is functional. What hampers the movement of the arm that is asleep is the missing contribution of dynamic control.

Summary

This section discussed distributed learning and remembering in the context of Lashley's (1929) work and biofunctional cognition (Iran-Nejad, 1980; Iran-Nejad and Ortony, 1984). The biofunctional model assumes that brain components are specialized subsystems and microsystems, as opposed to a patchwork of regions, centers, areas, nerve segments, or chemical particles. Brain microsystems are distributed, with regard to their physical location, in localized groups or in widely-scattered constellations. Lashley's concept of equipotentiality was analyzed biofunctionally in several different ways. We argued that equipotentiality and specialization in the nervous system are compatible. Indeed, although equipotentiality contrasts with specificity of memory traces, it is compatible with the specialization of brain subsystems and microsystems. Our analysis showed that Lashley did not certainly go too far, as some have claimed, in assuming equipotentiality. Factors that contribute to learning and remembering have their origin in different sources of control and learning processes and any act of learning occurs under the simultaneous influence of many factors. In this sense, learning and remembering are multisource phenomena.

Conclusions

This paper examined distributed learning and remembering (DLR) from the perspectives of parallel distributed processing (PDP) connectionism, the DLR foundational research, and a biofunctional model of mental functioning. A number of tentative conclusions may be drawn based on this analysis:

1. Parallel distributed processing connectionism and DLR foundational research are paradigmatically incompatible.

2. To our understanding, the only interpretation possible of the term *distributed* in the context of PDP connectionism is with reference to a one-to-

many relationship in which *one* refers to a mental macrounit and *many* refers to a localized group of mental microunits (intralevel group distribution). Intralevel group distribution may be contrasted with a similar interlevel definition of the term in which *one* refers to a mental unit and *many* refers to a localized group of brain microsystems (interlevel group distribution). Intralevel group distribution faces the insurmountable problem of reduction implied by the Gestalt idea that the whole is more than the sum of its parts. Interlevel group distribution does not face this problem.

3. Interlevel group distribution occurs within particular brain subsystems and provides the basis for generating image-like mental displays that are not meaningful by themselves. Meaningful mental phenomena are created and upheld by distributed constellations of brain microsystems whose members are scattered in various subsystems throughout the nervous system (interlevel constellation distribution). In this sense, the meaning of a within-subsystem mental display is the combination of mental displays created and upheld simultaneously in all other brain subsystems.

4. An important concept related to DLR is Lashley's (1929) notion of equipotentiality. Parallel distributed processing connectionists have found this notion problematic mainly because it runs counter to the postulation of specific traces which are inevitable in connectionism. Our analysis of this concept in terms of biofunctional cognition suggests that there is nothing wrong with Lashley's equipotentiality hypothesis.

5. In the biofunctional approach, the term *distributed* refers to the location of brain microsystems in the three-dimensional nervous system. There is, however, another related sense in which learning and remembering may be viewed as distributed. In this sense, the factors that contribute to learning and remembering are distributed in that they come from many different internal and external sources. Learning and remembering are not responses to a single external source of input. They are truly multisource phenomena.

The biofunctional approach to DLR explains Lashley's (1929) findings and is compatible with other foundational DLR research (John, 1967, 1972). The origins of the approach even go beyond the immediate boundaries of DLR literature and extend into the broader realm of nonassociative research representative of the work of Dewey (1896), Angell (1907), Head (1920), and Bartlett (1932). In this broader form, the approach promises to serve as a solid foundation for the study of how the brain creates and upholds the mind.

References

Anderson, J.A., and Hinton, G.E. (1981). Models of information processing in the brain. In G.E. Hinton and J.A. Anderson (Eds.), *Parallel models of associative memory* (pp. 9–48). Hillsdale, New Jersey: Erlbaum.

Anderson, J.R. (1983). *The architecture of cognition*. Cambridge, Massachusetts: Harvard University Press.

Angell, J.R. (1907). The province of functional psychology. *Psychological Review, 14*, 61–91.

Bartlett, F.C. (1932). *Remembering*. Cambridge, England: Cambridge University Press.

Bereiter, C. (1985). Toward a solution to the learning paradox. *Review of Educational Research, 55*, 201–226.

Bereiter, C. (1991). Implications of connectionism for thinking about rules. *Educational Researcher, 20*(3), 10–16.

Berlyne, D.E. (Ed.). (1974). *Studies in new experimental aesthetics*. New York: Wiley.

Bernstein, J. (1981, December 14). Profiles: AI, Marvin Minsky. *The New Yorker*, pp. 50–126.

Black, J.B., and Wilensky, R. (1979). An evaluation of story grammars. *Cognitive Science, 3*, 213–230.

Bransford, J.D., Nitsch, K.E., and Franks, J.J. (1977). Schooling and the facilitation of knowing. In R.C. Anderson, R.J. Spiro, and W.E. Montague (Eds.), *Schooling and the acquisition of knowledge* (pp. 31–55). Hillsdale, New Jersey: Erlbaum.

Broadbent, D. (1985). A question of levels: Comment on McClelland and Rumelhart. *Journal of Experimental Psychology: General, 114*, 189–192.

Chomsky, N. (1959). Review of Skinner's *Verbal Behavior*. *Language, 35*, 26–58.

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, Massachusetts: MIT Press.

Dewey, J. (1896). The reflex arc concept in psychology. *The Psychological Review, 3*(4), 357–370.

Diener, E., and Iran-Nejad, A. (1986). The relationship in experience between various types of affect. *Journal of Personality and Social Psychology, 50*, 1031–1038.

Head, H. (1920). *Studies in neurology*. London: Hodder and Stoughton.

Hebb, D.O. (1949). *The organization of behavior*. New York: Wiley.

Hinton, G.E., and Anderson, J.A. (Eds.). (1981). *Parallel models of associative memory*. Hillsdale, New Jersey: Erlbaum.

Hubel, D.H., and Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology, 195*, 215–243.

Hubel, D.H., and Wiesel, T.N. (1977). Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society of London, Series B, 198*, 1–59.

Iran-Nejad, A. (1980, February). *The schema: A structural or a functional pattern* (Tech. Rep. No. 159). Urbana, Illinois: University of Illinois, Center for the Study of Reading.

Iran-Nejad, A. (1987). Cognitive and affective causes of interest and liking. *Journal of Educational Psychology, 79*, 120–130.

Iran-Nejad, A. (1989a). Associative and nonassociative schema theories of learning. *Bulletin of Psychonomic Society, 27*, 1–4.

Iran-Nejad, A. (1989b). A nonassociative schema theory of cognitive incompatibility. *Bulletin of the Psychonomic Society, 27*, 429–432.

Iran-Nejad, A. (1989c). A nonconnectionist schema theory of understanding surprise-ending stories. *Discourse Processes, 12*, 127–148.

Iran-Nejad, A., and Chissom, B. (1992). Contributions of active and dynamic self regulation in learning. *Innovative Higher Education, 17*(2), 125–136.

Iran-Nejad, A., Clore, G. L., and Vondruska, R. J. (1984). Affect: A functional perspective. *The Journal of Mind and Behavior, 5*, 279–310.

Iran-Nejad, A., McKeachie, W.J., and Berliner, D.C. (1990). The multisource nature of learning: An introduction. *Review of Educational Research, 60*, 509–515.

Iran-Nejad, A., and Ortony, A. (1984). A biofunctional model of distributed mental content, mental structures, awareness, and attention. *The Journal of Mind and Behavior, 5*, 171–210.

Iran-Nejad, A., and Winsler, A. (2000). Bartlett's schema theory and modern accounts of learning and remembering. *The Journal of Mind and Behavior, 21*, 5–36.

Jenkins, J.J. (1974). Remember that old theory of memory? Well, forget it. *American Psychologist, 29*, 785–795.

John, E.R. (1967). *Mechanisms of memory*. New York: Academic Press.

John, E.R. (1972). Switchboard versus statistical theories of learning and memory. *Science, 177*, 850–864.

Kuhn, T.S. (1962). *The structure of scientific revolutions*. Chicago: The University of Chicago Press.

Lashley, K.S. (1915). *Notes on the nesting activities of the noddy and snooty terns* (Publication No. 211, pp. 62–83). Washington, DC: Carnegie Institute.

Lashley, K.S. (1929). *Brain mechanisms and intelligence*. Chicago: The University of Chicago Press.

Lashley, K.S. (1950). In search of the engram. In *Symposia of the Society for Experimental Biology (No: 4): Psychological mechanism in animal behavior* (pp. 454–482). New York: Academic Press.

Lashley, K.S. (1951). The problem of serial order in behavior. In L.A. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 112–136). Pasadena, California: California Institute of Technology.

Marr, D. (1982). *Vision*. San Francisco: Freeman.

McClelland, J.L., and Rumelhart, D.E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *The Psychological Review, 88*, 375–407.

McClelland, J.L., and Rumelhart, D.E. (1985a). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General, 114*, 159–188.

McClelland, J.L., and Rumelhart, D.E. (1985b). Levels indeed! A response to Broadbent. *Journal of Experimental Psychology: General, 114*, 193–197.

McClelland, J.L., and Rumelhart, D.E. (1986). A distributed model of human learning and memory. In J.L. McClelland, D.E. Rumelhart, and the PDP Research Group (Eds.), *Parallel distributed processing* (Volume 2, pp. 170–215). Cambridge, Massachusetts: MIT Press.

McClelland, J.L., Rumelhart, D.E., and Hinton, G.E. (1986). The appeal of parallel distributed processing. In J.L. McClelland, D.E. Rumelhart, and the PDP Research Group (Eds.), *Parallel distributed processing* (Volume 1, pp. 3–44). Cambridge, Massachusetts: MIT Press.

McClelland, J.L., Rumelhart, D.E., and the PDP Research Group (Eds.). (1986). *Parallel distributed processing*. Cambridge, Massachusetts: MIT Press.

McCulloch, W.S. (1965). Finality and form in nervous activity. In W.S. McCulloch (Ed.), *Embodiments of mind* (pp. 256–275). Cambridge, Massachusetts: MIT Press.

McCulloch, W.S., and Pitts, W. (1943). A logical calculus of the ideas imminent in nervous activity. *Bulletin of Mathematical Biophysics, 5*, 115–133.

Medin, D.L., and Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Reviews, 85*, 207–238.

Medin, D.L., and Schwanenflugel, P.J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 355–368.

Minsky, M., and Pappert, S. (1969). *Perceptions*. Cambridge, Massachusetts: MIT Press.

Murphy, G.L., and Medin, D.L. (1985). The role of theories in conceptual coherence. *The Psychological Review, 92*, 289–316.

Neisser, U. (1967). *Cognitive psychology*. New York: Appleton–Century–Crofts.

Oden, G.C. (1988, November). *Connectionism may be more different than you think but less different than you hope*. Paper presented at the twenty-ninth annual meeting of the Psychonomic Society, Chicago, Illinois.

Pinker, S., and Mehler, J. (Eds.). (1988). *Connections and symbols*. Cambridge, Massachusetts: MIT Press.

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 symbols* (pp. 73–193). Cambridge, Massachusetts: MIT Press.

Polanyi, M. (1958). *Personal knowledge*. Chicago: University of Chicago Press.

Prince, A., and Pinker, S. (1988). Rules and connections in human language. *TINS, 11*(5), 195–202.

Reece, P. (1987, January). Perceptrons and neural nets. *AI Expert*, pp. 50–57.
Rosenblatt, F. (1962). *Principles of neurodynamics*. New York: Spartan.
Rumelhart, D.E. (1975). Notes on a schema for stories. In D.G. Bobrow and A. Collins (Eds.), *Representation and understanding: Studies in cognitive science* (pp. 211–236). New York: Academic Press.
Rumelhart, D.E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and performance* (Volume 1, pp. 573–603). Hillsdale, New Jersey: Erlbaum.
Rumelhart, D.E. (1978). *Schemata: The building blocks of cognition* (CHIP Report No. 79). La Jolla, California: University of California, Center for Human Information Processing.
Rumelhart, D.E. (1980). Schemata: The building blocks of cognition. In R.J. Spiro, B.C. Bruce, and W.F. Brewer (Eds.), *Theoretical issues in reading comprehension: Perspectives from cognitive psychology, linguistics, artificial intelligence, and education* (pp. 33–58). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
Rumelhart, D.E., Hinton, G.E., and McClelland, J.L. (1986). A general framework for parallel distributed processing. In D.E. Rumelhart, J.L. McClelland, and the PDP Research Group (Eds.), *Parallel distributed processing* (Volume 1, pp. 45–76). Cambridge, Massachusetts: MIT Press.
Rumelhart, D.E., and McClelland, J.L. (1982). An interactive activation model of context effects in letter perception: Part 2. The contextual enhancement effect and some tests and extensions of the model. *The Psychological Review*, *89*, 60–94.
Rumelhart, D.E., and McClelland, J.L. (1986). PDP models and general issues in cognitive science. In D.E. Rumelhart, J.L. McClelland, and PDP Research Group (Eds.), *Parallel distributed processing* (Volume 1, pp. 110–146). Cambridge, Massachusetts: MIT Press.
Rumelhart, D.E., McClelland, J.L., and the PDP Research Group. (Eds.). (1986). *Parallel distributed processing*. Cambridge, Massachusetts: MIT Press.
Rumelhart, D.E., Smolensky, P., McClelland, J.L., and Hinton, G.E. (1986). Schemata and sequential thought processes in PDP models. In J.L. McClelland, D.E. Rumelhart, and the PDP Research Group (Eds.), *Parallel distributed processing* (Volume 2, pp. 7–57). Cambridge, Massachusetts: MIT Press.
Smolensky, P. (1988). On the proper treatment of connectionism. *The Behavioral and Brain Sciences*, *11*, 1–74.
Sperry, R.W. (1943). Visuomotor coordination in the newt (*triturus viridescent*) after regeneration of the optic nerve. *Journal of Comparative Neurology*, *79*, 33–55.
Sperry, R.W. (1966). Selective communication in nerve nets: Impulse specificity vs. connection specificity. In F.O. Schmitt and T. Melnechuk, (Eds.), *Neurosciences research symposium summaries* (Volume 1, pp. 213–219). Cambridge, Massachusetts: MIT Press.
Weiss, P. (1936). Selectivity controlling the central–peripheral relations in the nervous system. *Biological Review*, *11*, 494–531.
White, B.W., Saunders, F.A., Scadden, L., Bach-Y-Rita, P., and Collins, C.C. (1970). Seeing with the skin. *Perception and Psychophysics*, *7*, 23–27.
Wilensky, R. (1983). Story grammars versus story points. *The Behavioral and Brain Sciences*, 6, 579–623.
Wimsatt, W.C. (1976). Reductionism, levels of organization, and the mind–body problem. In G.G. Globus, G. Maxwell, and I. Savodnik (Eds.), *Consciousness and the brain* (pp. 205–267). New York: Plenum Press.
Winston, P.H. (1975). Learning structural descriptions from examples. In P.H. Winston (Ed.), *The psychology of computer vision* (pp. 157–209). New York: McGraw–Hill.