Future Directions

Parallel distributed processing has come a long way since we began our explorations. Four years ago, we had no clear idea how PDP models were going to overcome the limitations of the perceptron; how they could be related to higher level concepts such as the schema; or how they could be applied fruitfully to language processing. We think it is fair to say that the chapters in both volumes of this book document considerable progress on these and other problems. During this same period, work emerging from other research groups has also produced major advances. Yet, in spite of the progress, these explorations are not complete. The more we have learned, the more questions we have uncovered. Thus, we think of these volumes more as a progress report than as a final statement.

Even as the book is going to press, the work is continuing and expanding. Our own ideas are evolving, and many other researchers are beginning to join in these explorations, bringing new viewpoints and new issues with them. In the following few paragraphs we sketch what we see as the major emerging themes of the next several years of work—the most promising directions and the biggest challenges that lie ahead.

From where we stand, three general areas stand out as major foci of future work. These are:

- The application of parallel distributed processing to higher-level cognitive processes and sequential symbol processing.
- Further development of mechanisms of learning and a concommitant analysis of network architectures.
- The closer binding of PDP models with neuroscience.

We discuss each of these developments below.

Higher-Level Processes

We began our work on PDP models with applications that were more perceptual in nature and that focused more on those areas in which conventional sequential symbol-processing models have been least successful. We obviously believe this is only a first step. We believe that PDP models will provide useful insights into such higher-level cognitive processes as thinking and langauge processing. Chapters 14 and 19 offer our initial explorations of these areas. We do not mean to suggest that we believe these chapters provide all of the answers. Indeed, we have already begun to elaborate these ideas. More important, many other researchers have begun to appreciate the benefits of massively parallel processing in these areas, and are beginning to develop models that go beyond the first steps we have taken in these two chapters. Here we indicate a bit more of what we see as the general form these new directions will take.

Sequential symbol processing. It is becoming obvious to many others besides ourselves that parallel processing and sequential symbol processing must somehow occur in the same system. There appear to be three major schemes for integrating these two processing modalities. One of these is our approach to the problem: to develop PDP models in which sequential behavior is captured in the successive settlings of a parallel network or set of networks. Touretzky and Hinton (1985), have been working somewhat along the same lines, developing a PDP production system. They hope eventually to do variable binding and instantiation and, at the same time, preserve the benefits of parallel distributed processing in their symbol-processing machine.

Beyond this approach, there appear to be two other approaches to what might be called *massive parallelism in symbol processing.* The first of these, suggested by Norman in Chapter 26, is the idea of a hybrid system—a PDP net or set of nets, coupled with a symbol processing controller. A similar approach is taken by Schneider (1985) in his recent model of controlled and automatic processes in attention. A slight variant of this approach is to embed calls to parallel modules inside a conventional sequential-processing system.

We applaud this kind of approach since it strives to preserve the best of both worlds. It also captures the fact that the cognitive system is complex and consists of many parts—a point that no one would dispute. We do not take this road ourselves, in part because we see no reason to suppose that the mechanisms that control cognitive processing are not themselves constructed of the same underlying parallel hardware as the other aspects of the cognitive systems, and in part because we prefer to view the system not so much in terms of controlled and controlling modules, but in terms of more distributed forms of control.

The final approach is to modify other computational frameworks so that they incorporate many of the features of PDP models. This approach is exemplified by John R. Anderson's (1983) production system implementation of our interactive activation model, and Thibadeau, Just, and Carpenter's (1982) CAPS production system.

These developments are signs that other researchers, coming from other computational viewpoints, have begun to appreciate many of the benefits of interactive parallel processing. By building parallel capabilities directly into production systems, however (that is, by effectively allowing large numbers of productions to partially match at the same time), these models have some tendency of becoming top-heavy. It remains to be seen whether the direct computational benefits these models have (e.g., explicit, direct variable binding of many productions at the same time) will turn out to outweigh the extra computational overhead they introduce. Our own hope is that we can succeed in achieving models that capture the extent of human capabilities more naturally as emergent properties of parallel networks, but in the meantime, these massively parallel production systems provide another way to bring parallel distributed processing together with the sequential symbol-processing tradition in cognitive science.

Language processing. We began our work on PDP models with a strong interest in linguistic information processing. It was an attempt to deal with the problems we saw with more conventional accounts of language that contributed to our original attraction to PDP models. Some of the basic directions of our account of language are indicated in Chapters 18 and 19. We see the fleshing out of these sketches as central to our understanding of language acquisition, syntactic processing, and the interaction between syntactic and semantic processing.

Although it is not yet represented directly in our work, we believe that our approach will offer new insight into the role of metaphor and analogy in language. Chapter 19 begins to indicate that the notion of literal meaning, like the notions of grammaticality in syntax and regularity in morphology, will turn out to be a matter of degree and will best be explained in terms of coalitions of units interacting to produce the observed regularities. The notions of figurative and literal meanings will best be seen as coarse categories describing the nature of the meanings synthesized by PDP networks rather than fundamentally different meaning types arising from fundamentally different processes. Meanwhile, it is clear to everyone that language processing provides many challenges and many opportunities to test and elaborate our developing conception of how to build powerful processing machinery out of simple computational elements, while at the same time preserving the beneficial aspects of parallel distributed processing.

Learning and Architecture

Learning has always been a key issue in the development of network models. One of the most appealing features of PDP networks is that very simple, homogeneous learning procedures can be used to allow the networks to self-modify and adapt to their environments. Until very recently, these self-organizing systems have been of limited complexity. No multilayer learning algorithm has existed. Two such learning procedures have been developed in this book (Chapters 7 and 8) and others are under development elsewhere (cf. Barto & Anandan, 1985). These learning schemes are already generating a large amount of activity. With the development of the error propagation learning procedure, for example (Chapter 8), we can study the kinds of internal representations that evolve to support processing in particular problems and, in that way, come to understand better the problems under study as well as the nature of the networks required to solve them. Moreover, we can use these powerful learning procedures for genuine artificial intelligence—that is, as a tool for designing systems capable of discovering new ways of solving difficult problems.

Neuroscience

In spite of the fact that PDP models exploit brain-style processing, our own work has not yet moved far toward making direct contact with neuroscience. We believe that the opportunity is there, however, and that the future will bring more and more models specifically designed with data from neuroscience in mind. There are two natural areas of application—the application to data from people with brain pathologies and the application of PDP models to detailed modeling of neurophysiological and neuroanatomical data.

Neuropsychology. The study of cognitive deficits that arise from brain damage has focused on dissociations of function. If one patient shows one deficit, and another shows a different deficit, it is reasonable to infer that their different pathologies damaged different, separable components of the cognitive system. From this kind of data, one can try to characterize the separate components of the cognitive system and to identify the role each component plays in the functioning of the system as a whole. This approach has been successful in leading to several important insights about the macrostructure of the cognitive system. But there has been little work, until now, that focuses on the development of models that attempt to characterize the microstructure of the components in such as way as to provide an account of the way in which performance degrades with damage. Cognitive dysfunction is a matter of degree, so different patients can manifest widely different degrees of what appears to be a common deficit. Mild cases often show their deficits in subtle ways, and make relatively subtle errors—confusing similar concepts but not confusing distinct concepts. Some pathologies that produce highly diffuse damage (e.g., Alzheimer's disease) are progressive and increase gradually in severity as the disease process continues.

PDP models hold out considerable promise of providing detailed accounts of the way in which cognitive function degrades with damage. This point has been made by J. A. Anderson (1977; Anderson, Silverstein, Ritz, & Jones, 1977) and by Wood (1978), and models of aspects of neuropsychological phenomena can be found in two places in this book (Chapters 3 and 25). Much more can now be done, we think, to further the application of PDP models to neuropsychological phenomena. Equally important, neuropsychological data can, in some cases, help us make decisions about the macrostructure of the cognitive systems we construct out of PDP networks. Thus, there is hope that we will soon see real synergistic interactions between neuropsychology and parallel modeling.

Physiology and anatomy. Our primary focus has been on "cognitive" data, but the language of PDP is, by design, well suited as a medium for actual neural modeling. We expect that as these theoretical tools become more familiar and better understood, they will prove more and more useful for organizing data from neuroscience. One clear example of this lies in the paper by Gluck and Thompson (in press), in which PDP modeling techniques are applied to an analysis of conditioning in the simple nervous system of *Aplysia*, as studied by Kandel and his collaborators (Hawkins, Abrams, Carew, & Kandel, 1983; Kandel & Schwartz, 1982). Although we are a long way from connecting our more abstract networks with particular brain structures, we believe that the application of PDP models to cognitive and neural modeling enterprises will serve to both clarify the specifically neuroscience questions and lead us to a better understanding of the actual neural substructure underlying cognitive processes in the brain.

In addition to the neurophysiology of the brain, we expect that neuroanatomical facts will have an increasing impact on the way we build our PDP models. We have learned that the major determiner of what a particular system learns is largely a function of the possible connections. It is the anatomy that determines the possible connections. It seems likely that in subsystems in which the anatomy is relatively well understood, as for example in the visual system, that we can use the known anatomical structure in our models to constrain them and thereby explain why information is processed in the way that it is.

The four years we have spent working on these volumes have been exciting ones for all of us. We have learned much. It is our hope in publishing these volumes that they will communicate some of this sense of excitement, and that at least some of our readers will take up the problems and issues where we have left them, will help us correct some of our errors and misconceptions, and will help advance our understanding of the microstructure of cognition.

David E. Rumelhart
James L. McClelland