CSC: Classic Paper Review/Analysis

Title and Author

Title The Appeal of Parallel Distributed Processing

Author J. L. McClelland, D. E. Rumelhart, and G. E. Hinton

Summary/Hook

In this chapter, McClelland, Rumelhart, and Hinton introduce, describe, and explain Parallel Distributed Processing (PDP) models and their many applications in the field of cognitive science. PDP seeks to define and develop a processing system that more adequately models the many state-to-state transitions that comprise the complex internal structure of human cognition, primarily through a neural networking perspective, serving as an alternative argument to the more traditional symbolic/representational approach. Elements of artificial intelligence, psychology, linguistics, education and learning, and neuroscience are discussed in-depth, where the authors offer a variety of examples in support of the connectionist approach in contrast to the symbolic/representational approach in the modeling of cognition. Aspects of human cognition, such as natural language processing (and the interplay between syntax and semantics in knowledge and understanding), thought processes, learning and information processing, memory, vision, and perception are explored throughout this chapter through the application of connectionist and neural networking principles.

Knowledge Relating to the Cognitive Science Program Learning Outcomes

1. **Symbol Systems**

   However, it is important to bear in mind that most everyday situations cannot be rigidly assigned to just a single script. They generally involve an interplay between a number of different sources of information. Consider, for example, a child's birthday party at a restaurant. We know things about birthday parties, and we know things about restaurants, but we would not want to assume that we have explicit knowledge (at least, not in advance of our first restaurant birthday party) about the conjunction of the two. Yet we can imagine what such a party might be like. The fact that the party was being held in a restaurant would modify certain aspects of our expectations for birthday parties (we would not expect a game of Pin-the-Tail-on-the-Donkey, for example), while the fact that the event was a birthday party would inform our expectations for what would be ordered and who would pay the bill. Representations like scripts, frames, and schemata are useful structures for encoding knowledge, although we believe they only approximate the underlying structure of knowledge representation that emerges from the class of models we consider in this book, as explained in Chapter 14. Our main point here is that any theory that tries to account for human knowledge using script-like knowledge structures will have to allow them to interact with each other to capture the generative capacity of human understanding in novel situations. Achieving such interactions has been one of the greatest difficulties associated with implementing models that really think generatively using script- or frame-like representations.

2. Neural Networking

One reason for the appeal of PDP models is their obvious " physiological" flavor: They seem so much more closely tied to the physiology of the brain than are other kinds of information-processing models. The brain consists of a large number of highly interconnected elements (Figure 3) which apparently send very simple excitatory and inhibitory messages to each other and update their excitations on the basis of these simple messages. The properties of the units in many of the PDP models we will be exploring were inspired by basic properties of the neural hardware. In a later section of this book, we will examine in some detail the relation between PDP models and the brain. Though the appeal of POP 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 PDP models appeal to us for psychological and computational reasons. They 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.

3. Consciousness and Controversies

The process of human cognition, examined on a time scale of seconds and minutes, has a distinctly sequential character to it. Ideas. come, seem promising, and then are rejected; leads in the solution to a problem are taken up, then abandoned and replaced with new ideas. Though the process may not be discrete, it has a decidedly sequential character, with transitions from state-to-state occurring, say, two or three times a second. Clearly, any useful description of the overall organization of this sequential flow of thought will necessarily describe a sequence of states. But what is the internal structure of each of the states in the sequence, and how do they come about? Serious attempts to model even the simplest macrosteps of cognition-say, recognition of single words-require vast numbers of microsteps if they are implemented sequentially. As Feldman and Ballard (1982) have pointed out , the biological hardware is just too sluggish for sequential models of the microstructure to provide a plausible account , at least of the microstructure of human thought. And the time limitation only gets worse, not better, when sequential mechanisms try to take large numbers of constraints into account. Each additional constraint requires more time in a sequential machine, and, if the constraints are imprecise, the constraints can lead to a computational explosion. Yet people get faster, not slower, when they are able to exploit additional constraints. Parallel distributed processing models offer alternatives to serial models of the microstructure of cognition. They do not deny that there is a macrostructure, just as the study of subatomic particles does not deny the existence of interactions between atoms. What PDP models do is describe the internal structure of the larger units, just as subatomic physics describes the internal structure of the atoms that form the constituents of larger units of chemical structure.

4. Algorithms and Automata

Marr and Poggio (I976) began by explicitly representing the two views in two arrays, as human observers might in two different retinal images. They noted that corresponding black dots at different perceived distances from the observer will be offset from each other by different amounts in the two views. The job of the model is to determine which points correspond. This task is, of course , made difficult by the fact that there will be a very large number of spurious correspondences of individual dots. The goal of the mechanism, then, is to find those correspondences that represent real correspondences in depth and suppress those that represent spurious correspondences. To carry out this task, Marr and Poggio assigned a processing unit to each possible conjunction of a point in one image and a point in the other. Since the eyes are offset horizontally, the possible conjunctions occur at various offsets or disparities along the horizontal dimension. Thus , for each point in one eye, there was a set of processing units with one unit assigned to the conjunction of that point and the point at each horizontal offset from it in the other eye. Each processing unit received activation whenever both of the points the unit stood for contained dots. So far, then, units for both real and spurious correspondences would be equally activated. To allow the mechanism to find the right correspondences, they pointed out two general principles about the visual world: (a) Each point in each view generally corresponds to one and only one point in the other view, and (b) neighboring points in space tend to be at nearly the same depth and therefore at about the same disparity in the two images. While there are discontinuities at the edges of things, over most of a two dimensional view of the world there will be continuity. These principles are called the uniqueness and continuity constraints, respectively. Marr and Poggio incorporated these principles into the interconnections between the processing units. The uniqueness constraint was captured by inhibitory connections among the units that stand for alternative correspondences of the same dot. The continuity principle was captured by excitatory connections among the units that stand for similar offsets of adjacent dots.

5. Psychological Investigations

One very prominent feature of human memory is that it is content addressable. It seems fairly clear that we can access information in memory based on nearly any attribute of the representation we are trying to retrieve. Of course, some cues are much better than others. An attribute which is shared by a very large number of things we know about is not a very effective retrieval cue, since it does not accurately pick out a particular memory representation. But, several such cues, in conjunction can do the job. Thus, if we ask a friend who goes out with several women, "Who was that woman I saw you with?", he may not know which one we mean-but if we specify something else about her-say the color of her hair, what she was wearing (in so far as he remembers this at all), where we saw him with her-he will likely be able to hit upon the right one. It is, of course, possible to implement some kind of content addressability of memory on a standard computer in a variety of different ways. One way is to search sequentially, examining each memory in the system to find the memory or the set of memories which has the. particular content specified in the cue. An alternative, somewhat more efficient, scheme involves some form of indexing- keeping a list, for every content a memory might have, of which memories have that content. Such an indexing scheme can be made to work with error-free probes, but it will break down if there is an error in the specification of the retrieval cue. There are possible ways of recovering from such errors, but they lead to the kind of combinatorial explosions which plague this kind of computer implementation. But suppose that we imagine that each memory is represented by a unit which has mutually excitatory interactions with units standing for each of its properties. Then, whenever any property of the memory became active, the memory would tend to be activated, and whenever the memory became activated, all of its contents would tend to become activated. Such a scheme would automatically produce content addressability for us.

Though it would not be immune to errors , it would not be devastated by an error in the probe if the remaining properties specified the correct memory.