

Parallel Distributed Processing

Implications for Psychology and Neurobiology

Edited by

R. G. M. MORRIS

University of Edinburgh

CLARENDON PRESS · OXFORD

1989

Parallel distributed processing: implications for cognition and development¹

JAMES L. MCCLELLAND

Introduction

What kind of processing mechanism is the mind? Is it a sequential information processing machine, like the von Neumann computer? Or is it a massively parallel processor? The fact that human thought takes place in a device consisting of some tens of billions of neurons seems to support the parallel view. Yet until recently, there has been little attention to this fact among those who study the higher mental processes, and little convergence in the study of mind and brain.

Feldman (1981) has pointed out that the human brain places constraints on the methods that might be used to implement human thought. Neurons are relatively sluggish, noisy processing devices, compared to today's computers. Yet people can perceive a visual scene at a glance and recognize an object in about half a second. Feldman estimates that this leaves time for perhaps a hundred processing steps; but sequential algorithms for perception generally require hundreds of thousands. The facts imply that we exploit the brain's obvious capacity for parallel processing.

In view of this, researchers have begun to work toward theories of mental processes that rely on these parallel capabilities. These models are variously known as parallel distributed processing (PDP) models, neural models, or, perhaps most generally, connectionist models. Work is proceeding in several directions. Cognitive scientists seeking to provide a characterization of the nature of human thought have turned to building computational models in which a number of interconnected processors work in concert in performing some information processing task. Meanwhile, neuroscientists seeking to understand the functional properties of neural circuits are also building computational models, exploring the collective properties of ensembles of neuron-like processing units. These enterprises often have somewhat different goals; yet each informs and enriches the other, and each is pursued with the hope that someday these two directions of research will converge upon a shared understanding of brain and mind.

In this chapter, I take primarily a cognitive perspective. First, I review the

connectionist framework, stressing how basic aspects of cognition representation, processing, knowledge, and learning are captured in the connectionist framework. Next, the framework is applied to fundamental questions about the development of human thought, and some of the implications that the framework has for basic questions about cognition and development are illustrated.

The connectionist framework

The term 'connectionist models' was introduced by Feldman (1981; Feldman and Ballard 1982). In these papers, the term was used to refer to a class of models that compute by way of connections among simple processing units. Another phrase often used to describe some connectionist models is *parallel distributed processing* or PDP models (Rumelhart, McClelland, and the PDP Research Group 1986; McClelland, Rumelhart, and the PDP Research Group 1986). PDP models are instances of connectionist models that stress the notion that processing activity results from the processing interactions occurring among rather large numbers of processing units.

In this article I intend the phrase 'the connectionist framework' to encompass all kinds of connectionist models. The framework may be thought of as providing a set of general assumptions about basic aspects of information processing, and a set of soft constraints on the range of specific assumptions that might be made. In what follows I consider each of several aspects of an information processing system. I describe the general assumptions connectionist models make about these aspects and I characterize some of the specific assumptions that might be made. The presentation draws heavily on Rumelhart, Hinton, and McClelland (1986), which can be consulted for further details.

Primitives and their organization

Like all cognitive models, connectionist models must propose some building blocks and some organization of these building blocks. In connectionist models, the primitives are *units* and *connections*. Units are simple processing devices which take on activation values based on a weighted sum of their inputs from the environment and from other units. Connections provide the medium whereby the units interact with each other; they are weighted, and the weights may be positive or negative, so that a particular input will tend to excite or inhibit the unit that receives it, depending on the sign of the weight (we shall return to these matters when we consider the dynamics of processing below).

Any particular connectionist model will make assumptions about the

number of units, their pattern of connectivity to other units, and their interactions with the environment. These assumptions define the architecture of a connectionist model. The set of units and their connections is typically called a *network*.

It should be noted that a very wide variety of architectures is possible. Two are shown in Figs 2.1 and 2.2. One of these, in Fig. 2.1 (from the distributed




Fig. 2.1 A fully connected autoassociator network, with connections from each unit to every other unit. Each unit receives input from outside the network, and sends output outside the network. All connections may be either positive or negative in this very general formulation. (From J. L. McClelland and D. E. Rumelhart, 1985, 'Distributed memory and the representation of general and specific information', *Journal of Experimental Psychology: General*, 114, 162. Copyright 1985 by the American Psychological Association. Reprinted by Permission.)

model of memory examined by McClelland and Rumelhart (1985) shows a set of completely interconnected units, each receiving input from the environment, and each projecting back to the environment. In some sense, the network in this diagram is the most general possible connectionist architecture, in that all others involve restrictions of this general case. For example, some units may receive no input from the environment; some may send no output outside the net; and some of the interconnections among units in the network may be deleted. There may, furthermore, be restrictions on the values of some of the connections. In the general case, each may be positive or negative, but the architecture may prescribe, for example, that a certain group of units have mutually inhibitory connections of fixed strength.

Figure 2.2 gives an example of a more restricted architecture, from the interactive activation model of visual word recognition (McClelland and Rumelhart 1981). In this model, units stand for hypotheses about displays of letter strings at each of three levels of description: a feature level, a letter level, and a word level. There are excitatory connections (in both directions)




Fig. 2.2 A sketch of the network used in the interactive activation model of visual word recognition (McClelland and Rumelhart 1981). Units within the same rectangle stand for incompatible alternative hypotheses about an input pattern, and are mutually inhibitory. Bi-directional excitatory connections between levels are indicated for one word and its constituents. (From J. L. McClelland, 1985, 'Putting knowledge in its place: a scheme for programming parallel processing structures on the fly', *Cognitive Science*, 9, 115. Copyright 1985 by Ablex Publishing. Reprinted by permission.)

between mutually consistent units on adjacent levels, and inhibitory connections between mutually inconsistent units within the same level. Thus the unit for T in the first letter position excites and is excited by the units for features of the letter T, as well as the units for words that begin with T. This unit also inhibits, and is inhibited by, units for other letters in the same letter position.

Active representation

Representations in connectionist models are patterns of activation over the units in the network. In some ways, these kinds of patterns are similar to representations in other frameworks; after all, representations in a computer are ultimately patterns of 0s and 1s. There are differences, however. For one thing it is quite natural for connectionist representations to be graded, in the sense that each unit's activation need not be one of two binary values. In some models, activations are restricted to binary or some other number of discrete values, but more typically each unit may take on a continuous activation value between some maximum and minimum. A more important difference is this: connectionist representations are truly active, in the sense that *they give rise to further processing activity directly*, without any need for a central processor or production-matching-and-application mechanism that examines them and takes action on the basis of the results of this examination.

Models differ in terms of the extent to which individual processing units can be identified with particular conceptual objects, such as letters, words, concepts, etc. The models illustrated in Figs 2.1 and 2.2 represent endpoints on a continuum. In the distributed model of memory, each conceptual object is thought of as a pattern of activation over a number of simple processing units. In the interactive activation model of word perception, on the other hand, each unit stands for a primitive conceptual object, such as a letter, a word, or a distinct visual feature. A large number of models lie between these two extremes (see Hinton, McClelland, and Rumelhart 1986, and Feldman 1988, for general discussions of the issue of distributed representation).

Processing

Processing in connectionist models occurs through the evolution of patterns of activation over time. This process is governed by assumptions about the exact way in which the activations of units are updated, as a function of their inputs. Updating can be synchronous (all units updated simultaneously) or asynchronous (units updated in random order). Updating generally occurs as follows. First, a *net input* is computed for each unit to be updated. The net input is the sum of the activations of all of the units that project to it, with each contributing activation weighted by the weight on the connection from the contributing unit to the receiving unit.² The net input may also include a *bias* term associated with the unit, as well as a term for input arising from outside the network. Thus for unit i , the net input is given by:

$$net_i = \sum_j w_{ij} a_j + bias_i + input_i$$

Here, j runs over all the units with connections projecting to unit i . The net input can then be used to set the new activation of the unit according to some monotonic but non-linear function such as the one shown in Fig. 2.3. Alternatively, the net input can be used to set the activation of the unit probabilistically to one of two discrete values (usually 1 or 0). Another possibility is that the net input may act as a force, tending to drive the activation of the unit up or down a small amount in each time step.³

It is typical to use some form of non-linear activation function, so that the activation of a unit is not simply set equal to the net input or some weighted average of the net input and the previous activation of the unit. Non-linearities are typically necessary for two reasons. 1. Linear networks are subject to explosive growth of activation due to positive feedback loops unless the weights are severely constrained (see Shrager, Hogg, and Huberman 1987). 2. Many computations require a layer of non-linear units between input and output. Without non-linearities, multiple layers of units add no additional




Fig. 2.3 The logistic function, a smooth non-linear function that is frequently used in relating activations of units to their net inputs. This function is often used to set the activation of a unit to a value between 0 and 1, or to set the probability of the unit to 1 or 0 probabilistically, with the probability determined by the value of the function.

computational power over that offered by a single layer (see Rumelhart, Hinton, and McClelland 1986, for further explanation).

Knowledge

Crucial to the very idea of cognition is the notion that information processing is guided by knowledge. We recognize the word THE as a definite article because of knowledge we have about the relation between letter strings and linguistic forms. We infer that a spoon may have been used if we hear 'the man stirred the coffee' because of knowledge we have about the kinds of instruments that are used for stirring. In many models, these kinds of knowledge would be stored in tables. For example, information about THE would be stored in a table called a lexicon, listing correspondences of letter strings and the linguistic objects they represent.

In connectionist models, knowledge is stored in the connections among the processing units. This assumption works together with the assumptions that connectionist models make about representations. An active representation on a set of units, together with the knowledge stored in connections, will give rise to new patterns of activation on the same or on other units.

Typically in connectionist models, connection strengths are real-valued. In models whose connections are set by assumption, it is typical to assume homogeneity of connection strengths as much as possible, to avoid excessive

degrees of freedom. In models that learn, however, connection strengths are typically allowed to take on whatever values the learning process gives them; parsimony arises from the use of a homogeneous principle of learning.

Learning

If knowledge is in the connection weights, learning must occur through the adjustment of these weights. This weight adjustment process is assumed to occur as a by-product of processing activity. Some knowledge can in fact be built into connectionist models, in the form of initial connection strengths, before there has been any learning, but it is common to explore the limits of what can be acquired through connection strength adjustment with minimal pre-wiring. The initial architecture of the network serves to impose constraints on the learning process; these can in many cases greatly facilitate learning and generalization, if these constraints are appropriate to the problem the network is given to learn.

A wide variety of 'learning rules' for tuning connections has been proposed. A recent review is provided by Hinton (in press). Generally, these rules state that the adjustment that is made to each connection should be based on the product of a 'pre-synaptic' term, associated with the unit that is sending input through the connection, and a 'post-synaptic' term, associated with the unit that is receiving input through the connection. For example, the *Hebb rule*, as used by Anderson (1977), makes the change in the strength of a connection proportional to the product of the activation of the sending unit and the receiving unit.⁴ Learning through connection strength adjustment is very different from learning processes in most other types of models. It is governed by simple mathematical expressions, and results in knowledge that is completely implicit, in that it is embedded inextricably in the machinery of processing, and is completely inaccessible to introspection or report. However, it should be noted, that while the connection changes themselves are not accessible, the patterns of activation whose construction they make possible can be accessible to other parts of the processing system.

The environment

Though it has been implicit in what I have said already, there is another aspect of connectionist models that deserves comment, namely their *environment*. The environment consists of an ensemble of possible patterns that might be presented to the network. In most cases, these patterns are thought of as separate events; each one presented when the network is in a resting state, then left on until processing is complete. However, input patterns can have a richer

temporal structure, or course; each event may consist of a sequence of events, or of a graded progression of input activations.

For networks with fixed connections, the environment simply defines the domain of inputs on which the network might be tested. For networks in which the connections are adjusted as a result of processing experience, however, the environment plays a crucial role in determining exactly what is learned. Thus models that aim to capture aspects of cognitive development through connectionist learning include among their assumptions a specification of the details of the experience that gives rise to the resulting developmental sequence. In many cases, these assumptions play a major role in determining the success or failure of the modelling effort.

The spirit of the thing

The connectionist framework is cast, not as a list of specific detailed assumptions, but as a set of *general principles* and some guidelines that provide weak constraints on the range of variants that fall within the scope of these principles. Indeed, as Rumelhart, Hinton, and McClelland (1986) noted, it is possible to build a von Neumann computer out of connectionist primitives, if they are organized in accordance with the von Neumann architecture. It thus becomes important to focus on the spirit of the connectionist framework. Generally, connectionist models of cognitive processes have been constructed expressly to exploit the capability for parallelism inherent in the approach, to make use of the graded capabilities of patterns of activation, and to capture the incremental nature of human learning in many tasks through the adjustment of connection strengths based on signals arising in the course of processing.

The microstructure of cognition

Finally, it is worth pointing out that the connectionist framework is not incompatible with other levels of description in cognitive science. Thus, there is nothing inconsistent with connectionist models in the claim that a cognitive system may traverse a sequence of states in a temporally extended cognitive task such as solving an arithmetic problem. According to the connectionist approach one would tend to view each such step in the process of solving the problem as a new state of the processing network. Indeed, Rumelhart *et al.* (1986) describe a network that performs a mental tic-tac-toe simulation, settling into a sequence of states representing the results of the successive mentally simulated moves made by each player.

There are important differences between conventional and connectionist models of sequential behaviour. In connectionist models, the states need not

be so discrete as they generally are in other models (Rumelhart and Norman 1982; Jordan 1986; Smolensky 1986). Furthermore, the powerful constraint-satisfaction characteristics inherent in the connectionist framework are not typically exploited by conventional models of sequential processing. The idea that each step is a sequential process involves a massively parallel constraint-satisfaction process seems like a promising starting place for a new way of thinking about the macrostructure of cognition.

The point that connectionist models characterize the microstructure of cognition applies not only in respect of time, but also in respect of the structure of the processing system and in respect of the description of the computational operations that the system is performing. Structurally, a processing system may consist of many parts, and for some purposes it may be adequate to describe its structure in terms of these parts and the flow of information between them. Computationally, too, it may often be useful and illuminating to describe what a part of such a system computes without referring specifically to the role in this computation that is played by the specific units and connections. The claim is, though, that it will be necessary to delve more deeply than this to provide a full description of the mechanisms of cognition.

Are connectionist models mere implementations?

In allowing that there may be a macrostructure to thought, connectionists may seem to suggest that their models merely describe the implementation details of a processing system that would be best characterized more abstractly. However, we simply do not know exactly what level of description is the appropriate one for characterizing many behavioural phenomena. Many who have turned to connectionist models have done so because these models have seemed to provide exactly the right level of description for characterizing certain kinds of cognitive processes. Just where the bounds of usefulness of the connectionist framework may lie seems at this point to be one of the very open questions. Since there is little in cognitive psychology that we understand perfectly at present, we are not in a position to say which aspects of cognition might be explainable without recourse to a model of the microstructure.

Connectionist models and cognitive development

In the preceding part of this chapter, I have tried to give an overview of the connectionist framework for cognitive modelling. Here, I consider the question: Does the connectionist framework have any implications for the answers that we give to basic questions about human cognition? I will argue

that it does. The questions are ones that arise within the field of cognitive development; they are motivated by dramatic behavioural phenomena. Several different kinds of answers have been given to these questions. We will see how the connectionist framework opens them anew and suggests what may be different answers in many cases.

The phenomena

The field of cognitive development is replete with examples of dramatic changes in children's thinking as they grow older. Here I give three examples: (1) failures of conservation and compensation; (2) progressive differentiation of knowledge about different kinds of things; and (3) U-shaped learning curves in language acquisition.

Failures of conservation and compensation

Perhaps the best-known phenomena in cognitive development are the dramatic failures of conservation that Piaget has reported in a wide range of different domains. One domain is the domain of liquid quantity. A child of 3 years is shown two glasses of water. The glasses are the same, and each contains the same amount of water, and the child sees that the amount is the same. But when the contents of one of the glasses is poured into a wider container, the child will say that there is less liquid in the wider container.

It is typical to say that this answer given by the young child reflects a failure to recognize two things: (1) that quantity is conserved under the transformation of pouring from one container to another; and (2) that greater width can compensate for less height. Many tasks are specifically designed to tap into the child's ability to cope with these kinds of compensation relations between variables.

One such task developed by Inhelder and Piaget (1958), the so-called *balance-beam task*, is illustrated in Fig. 2.4. In this task, the child is shown a




Fig. 2.4 Balance beam of the kind first used by Inhelder and Piaget (1958), and later used extensively by Siegler (1976, 1981; Siegler and Klahr 1982). (Reprinted from Siegler 1976, Fig. 1, with permission.)

balance beam with pegs placed at evenly spaced intervals to left and right of a fulcrum. On one peg on the left are several weights; on one peg on the right are several weights. The beam is immobilized, and the child is asked to judge which side will go down, or whether they will balance. We will have occasion to examine performance in this task at length below; for now it suffices to note that young children (up to about 6 or 7 years in this case) typically respond as if the distance from the fulcrum was completely irrelevant. They will say the beam should balance if the weight is the same on both sides, regardless of distance. Otherwise they say the side with the greater weight will go down. These children, then, appear to miss the fact that lesser weight can be compensated for by greater distance. Typically by the age of 11 years or so children have some appreciation for this trade off; the details of the developmental progression are quite interesting, as we shall see below.


Progressive differentiation of ontological categories

Other researchers, studying different domains, have noticed other kinds of developmental progressions. Keil (1979) studied children's judgements about whether you could say things like 'A rabbit is an hour long'. He supposed such judgements tapped children's knowledge about different kinds of things. In these judgements, Keil was interested not in whether the child saw a sentence as true or false, but in whether the child felt that one could make certain kinds of predictions (e.g. that something is an hour long) when the something is a member of a certain 'ontological category' (e.g. living thing). Keil found that children were much more permissive in their acceptance of statements than adults were, but their permissiveness was not simply random. Rather, children would accept statements that overextended predicates to categories near the ones that they typically apply to, but would not extend them further. Thus some children will accept predications like 'The rock is asleep', but not 'The rock is an hour long'. It was as though children's knowledge of what predicates apply to particular categories becomes progressively more and more differentiated, as illustrated in Fig. 2.5.


U-Shaped learning curves in language development

Early on, children often get certain kinds of linguistic constructions correct which they later get wrong; only much later do they recover their former correct performance. One example is the passive construction, applied to semantically biased materials, such as 'The man was bitten by the dog'. (See Bever, 1970, for a discussion of the development of the use of the passive construction.) Early in development, children correctly interpret such sentences; they appear to be using information about what roles the different


Kindergarten



Second Grade



Fourth Grade



Sixth Grade




Fig. 2.5 Four different 'predictability trees' illustrating the progressive differentiation of concepts as a function of age. Terms in capitals at internal nodes in the trees represent predicates, and terms in lower case at terminal nodes in the trees represent concepts that are spanned by all the predicates written on nodes that dominate the terminal. A predicate spans a concept if the child reports that it is not silly to apply either the predicate or its negation or both to the concept. Thus the first tree indicates that the child will accept 'The man is (not) alive'; and 'The chair is (not) tall', but will not accept 'The chair is (not) alive'. Parentheses indicate uncertainty about the application of a predicate. (Redrawn from Keil, 1979, Appendix C, with permission.)

nouns typically play in the action described by the verb, since they tend to be correct only when the correct interpretation assigns the nouns to their typical roles. At an older age, children respond differently to such sentences, treating the first noun-phrase as the subject; semantic constraints are over-ridden, and there is a tendency to interpret 'The man was bitten by the dog' as meaning 'The man bit the dog'. Finally, children interpret the sentence correctly again, but for a different reason. It would appear that they now understand passives in general, since at this stage they can also interpret semantically neutral and even reverse-biased sentences (such as 'The dog was bitten by the man') correctly.

The questions

The phenomena reviewed above raise basic questions about cognitive development. Three of these questions are:

1. Are these different phenomena simply unrelated facts about development in different domains?
2. Are there principles that all of these phenomena exemplify?
3. If there are principles, are they domain specific, or are they general principles about development?

Different kinds of developmental theorists have answered such questions in very different ways. To Piaget, each failure of compensation or conservation reflected a single common developmental stage; the phenomena were intrinsically related by the characteristics of the stage, and these characteristics provided the basis for explanation.

Others have taken a very different approach. Keil (1979), following Chomsky's analogous argument for language, argues for *domain specific principles* of development. His view is that each cognitive domain has its own laws that provide constraints on what can be learned. These constraints limit the hypotheses that the child can entertain, thereby making it dramatically easier for the child to acquire adult abilities in the face of the impoverished information that is provided by experience with the world.

The main thrust of the remainder of this chapter is to argue that recent developments in connectionist learning procedures suggest a dramatic alternative to these kinds of views. The alternative is simply the hypothesis that these diverse developmental phenomena all reflect the operation of a single basic learning principle, operating in different tasks and different parts of the cognitive system.

The learning principle

The principle can be stated in fairly abstract terms as follows:

Adjust the parameters of the mind in proportion to the extent to which their adjustment can produce a reduction in the discrepancy between expected and observed events.

This principle is not new. It might well be seen as capturing the residue of Piaget's accommodation process, in that accommodation involves an adjustment of mental structures in response to discrepancies. (See Flavell, 1963, for a discussion of Piaget's theory.) It is also very similar to the principle that governs learning in the Rescorla-Wagner model of classical conditioning (Rescorla and Wagner 1972). What is new is that there exists a learning procedure for multilayer connectionist networks that implements this principle. Here, the parameters of the mind are the connections among the

units in the network, and the procedure is the back-propagation procedure of Rumelhart, Hinton, and Williams (1986; see Hinton, this volume).

The learning principle lies at the heart of a number of connectionist models that learn how to do various different kinds of information processing tasks, and that have applications to phenomena in cognitive and/or language development. Perhaps the simplest such model is the past-tense model of Rumelhart and McClelland (1986). The development of that model pre-dated the discovery of the back-propagation learning procedure, thereby forcing certain simplifications for the sake of developing an illustration of the basic point that lawful behaviour might emerge from the application of a simple principle of learning to a connectionist network. Subsequent models have used back-propagation to overcome some of these limitations. Included in this class are NETalk (Sejnowski and Rosenberg 1987) and a more recent model of word reading (Seidenberg, Patterson, and McClelland, this volume). The present effort grew out of two observations of similarities between the developmental courses seen in models embodying this principle, and the courses of development seen in children: First, the course of learning in a recent model of concept learning by Rumelhart (in preparation) is similar to aspects of the progressive differentiation of concepts reflected in Keil's (1979) studies of predictability. Second, the course of learning in a recent model of sentence comprehension by St John and McClelland (1988) mirrors aspects of the progression from reliance on semantic constraints, to reliance on word order, to, finally, reliance on complex syntactic patterning such as the passive voice. I do not mean to claim that the models in question are fully adequate models of the developmental progression in either case; I only claim that they seemed suggestive: they raised the possibility that part of the explanation of these and other developmental phenomena might be found in the operation of the learning principle as it adjusts connection strengths in a network subjected to patterns arising in its environment.

The remainder of this chapter presents an experiment assessing the applicability of this conjecture to another developmental phenomenon, namely the acquisition of the ability to take both weight and distance into account in the balance beam task described above. The task has been studied extensively by Siegler and his colleagues (Siegler 1976, 1981; Siegler and Klahr 1982), and quite a bit is known about it. I will first review the developmental findings. Then I will describe a connectionist model that captures these phenomena by applying the learning principle stated above.

Development of judgements of balance

In an important monograph, Siegler (1981) studied children's performance in the balance beam task and three other tasks in which two cues had to be taken

into account for correct performance. In all cases, as in the balance beam task, the correct procedure requires multiplication. For example, in the balance beam task, to determine which side will go down, one must multiply the amount of weight on a given side of the beam by the distance of that weight from the fulcrum. The side with the greater product will go down; when the products are the same, the beam will balance.

Siegler studied children in several age groups, as well as young adults. Each child was asked to judge 24 balance problems. In each case, the beam was immobilized so that there was no feedback. The 24 problems could be divided into four of each of six types:

1. *Balance*. In this class of problem, the weight is the same on both sides of the beam and the weight is the same distance from the fulcrum on both sides.
2. *Weight*. In these problems, the weights differ but distance from the fulcrum is the same on both sides.
3. *Distance*. Here the weight is the same on both sides, but the distance from the fulcrum differs.
4. *Conflict*. Here both weight and distance differ and are in conflict, in that the weight is greater on one side but the distance from the fulcrum is greater on the other. There are three types of conflict problems:
 - (a) *conflict-weight*. In these cases, the side with the greater weight has the greater torque (i.e. the greater value of the product of weight times distance).
 - (b) *conflict-distance*. In these cases, the side with the greater distance has the greater torque.
 - (c) *conflict-balance*. Here the torques are the same on both sides.

Siegler's analysis of children's performance assumed that children use rule-governed procedures. Four such procedures or *rules* as Siegler called them are shown in Fig. 2.6. Each of these rules corresponds to a distinct pattern of performance over the six problem types. For example, children using Rule 1 should say the side with the greater weight will go down in weight problems and in all three types of conflict problems. They should think the beam will balance on balance problems and distance problems. In general, the mapping from the rules to expected performance is extremely straightforward. The only point that needs explication is the instruction *muddle through* when weight and distance conflict in Rule 3. In practice it is assumed to mean 'guess randomly among the alternatives', so that $\frac{1}{3}$ of the responses would be left-side-down; $\frac{1}{3}$ right-side-down, and $\frac{1}{3}$ balance.

Siegler compared the performance of each child tested with each rule, and counted discrepancies from predicted performance based on the rule. Children who scored less than four discrepancies from a given rule were scored as using that rule.

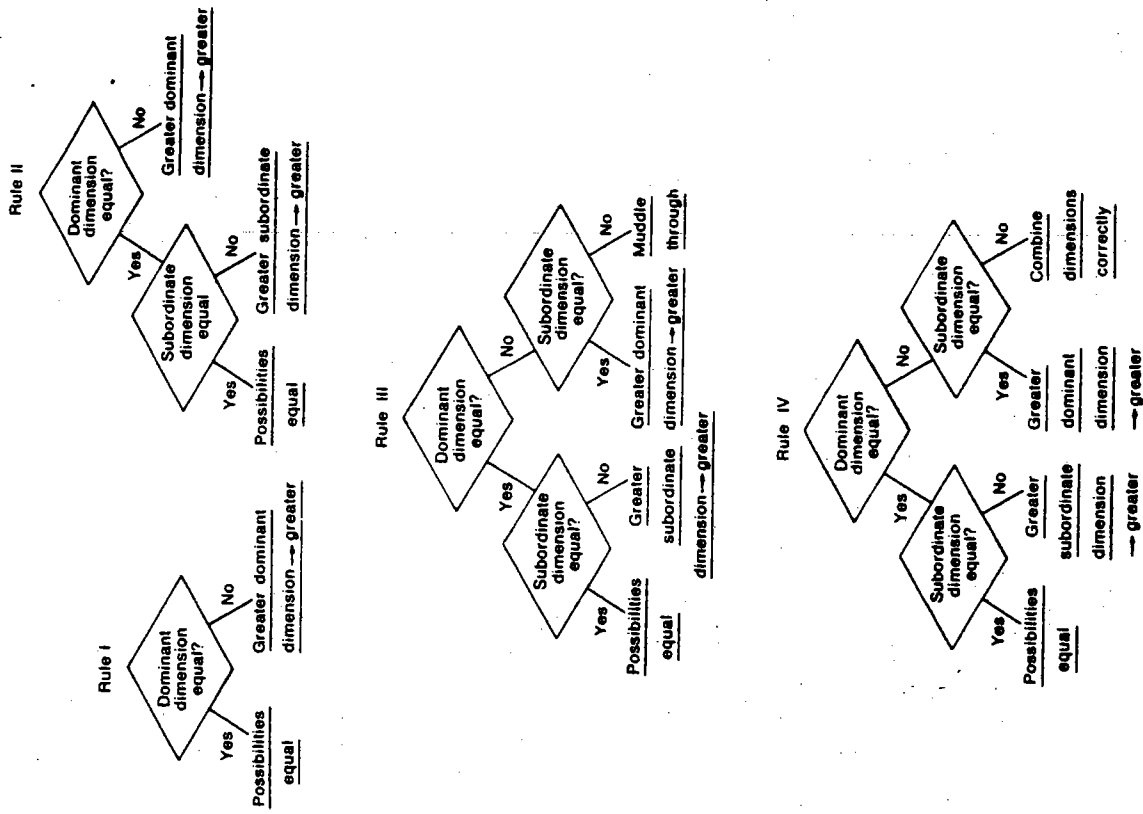


Fig. 2.6 Siegler's (1976, 1981) four 'rules' for answering balance beam problems. Each rule is in fact a full procedure, rather than a single rule. (Reprinted with permission from Siegler (1981), Fig. 1.)

For our purposes, there are four basic findings that emerge from Siegler's analysis:

1. *Lawful behaviour.* In general, performance of children over the age of 5 years is extremely regular in the balance beam task. Overall, about 90 per cent of children tested conform to one of the four rules.
2. *Developmental progression.* As children get older, they appear to progress through the use of the different rules. The progression from Rule 1 to Rule 3 can be thought of as a progression in which at first the weight cue is relied on exclusively, while at the end distance and weight are both taken into account. In between (Rule 2), distance is taken into account only if it does not conflict with the weight cue. Children aged 5-7 years typically use Rule 1, and college students typically use Rules 3 or 4. Many college students do not have explicit knowledge of the torque principle. Children younger than age 5 years tend not to be scorable strictly in terms of one of the rules; however, they appear to show an increasing tendency to behave in accordance with Rule 1.

3. *Generality.* The same four rules appear to be adequate to characterize performance in all three of the domains that Siegler studied. Though the developmental progression was not identical across cases, there was in all cases a trend from simpler to more complex rules with development.
4. *Lack of correlation between domains.* Even though children seem to progress through the same rules in different domains, they do not do so in lock-step; the correlation across domains is low, particularly in terms of the higher-numbered rules, so that children who are showing Rule 3 behaviour in one task may be showing Rule 1 behaviour in another and Rule 4 in a third.

The simulation model

The model I have developed to capture Siegler's findings is sketched in Fig. 2.7.⁵ Of course, the model is a drastic oversimplification of the human mind and of the task; but as we shall see it allows us to capture the essence of Siegler's findings, and to see them emerge from the operation of the learning principle described above.

The model consists of a set of input units, to which balance problems can be presented as patterns of inputs; a set of output units, over which the answer to each problem can be represented; and a set of hidden units, between the input and the output. Connections run from input units to hidden units and from hidden units to output units.

The input units can be divided into two groups of 10. One group is used to




Fig. 2.7 The network used in the simulation of the development of performance in the balance beam task.

represent information about weight and the other is used to represent information about distance. In each case I have chosen to use an input representation that imposes as little structure as possible on the input patterns. Each possible value of weight or distance from the fulcrum is assigned a separate unit. The ordering of values from low to high is not given in this representation; the network will have to learn this ordering. For the convenience of the reader, the units are arranged in rows according to which side of the beam they are from, and within each row they are arranged from left to right in order of increasing weight or distance from the fulcrum; but this ordering is unknown to the model before it is trained, as we shall see.

Though the two dimensions are not intrinsically structured for the model, the design of the network does impose a separate analysis of each dimension. This separation turns out to be critical; I will consider the implications of this architectural simplification below. The separation is implemented as follows: there are separate pairs of hidden units for each dimension. Two hidden units receive input from the weight input units and two receive input from the distance input units.

Each of the four hidden units projects to each of the two output units. The left output unit can be thought of as a 'left-side-down' unit, and the right one as a 'right-side-down unit'. Thus a correct network for the task would turn on the output unit corresponding to the side with the greater torque, and would turn off the unit for the other side. For balance problems, I assume that the network should turn both units on half-way. Note that this coding of the output

patterns does tell the network that balance is between left-side-down and right-side-down.

Processing

Balance problems of the kind studied by Siegler can be processed by the network by simply turning on (i.e. setting to 1) the input units corresponding to a particular problem and turning off (i.e. setting to 0) all other input units. The input from the problem illustrated in Fig. 2.7 is shown by using black to indicate those input units whose activations are 1.0, and white for the units whose activations are 0.

The inputs are propagated forward to the hidden units. Each hidden unit simply computes a net input:

$$net_i = \sum_j w_{ij} a_j + bias_i$$

Here j ranges over the input units. Each hidden unit then sets its activation according to the logistic function:

$$a_i = \frac{1}{1 + e^{-net_i}}$$

In these equations, w_{ij} is the strength of the connection to hidden unit i from input unit j , a_j is the activation of input unit j , and $bias_i$ is the modifiable bias of hidden unit i . This bias is equivalent to a weight to unit i from a special unit that is always on.

Once activations of the hidden units are determined, the activations of the output units are determined by the same procedure. That is, the net input to each output unit is determined based on the activations of the hidden units, the weights from the hidden units to the output units, and the biases of the output units. Then the activations of the output units are determined using the logistic function.

Responses

The activations of the output units are real numbers between 0 and 1; to relate the model's performance to the balance beam task, these real-valued outputs must be translated into discrete responses. I used the following simple translation: if the activation of one output unit exceeded the activation of the other by 0.333, I took the answer to be 'more active side down'. Otherwise, the answer was assumed to be 'both sides equal'.

Learning

Before training begins, the strengths of these connections from input to hidden units and from hidden to output units are initialized to random values uniformly distributed between +0.5 and -0.5. In this state, inputs lead to

random patterns of activity over both the hidden and output units. The activations of the output units fluctuate approximately randomly between about 0.4 and 0.6 for different input patterns. The network comes to respond correctly only as a result of training. Conceptually, training is thought of as occurring as a result of a series of experiences in which the network is shown a balance problem as input; computes activations of output patterns based on its existing connection weights; and is then shown the correct answer. The signal that drives learning is the difference between the obtained activation of each output unit and the correct or target activation for that unit. The back-propagation procedure of Rumelhart, Hinton, and Williams (1986) is then used to determine how each connection strength in the network should be adjusted to reduce these differences. The procedure is described in Hinton's chapter in this volume, and it would be redundant to describe it here. Suffice it to say that it exactly implements the learning principle stated above, and restated here in network terminology:

Adjust each weight in the network in proportion to the extent to which its adjustment can produce a reduction in the discrepancy between the expected event and the observed event, in the present context.

Here the 'expected event' is the pattern of activation over the output units that is computed by the network; the observed event is the pattern of activation that the environment indicates these units have; and the present context is the pattern of activation over the input units. Note that the direction of change to a connection (positive or negative) is simply the direction that tends to reduce the discrepancy between computed output and the correct or target output.

Environment

As I pointed out at the beginning of this chapter, the environment in which a network learns plays a very strong role in determining what it learns, and particularly the developmental course of learning. The simulations reported here were based on the assumption that the environment for learning about balance problems consists of experiences that vary more frequently on the weight dimension than they do on the distance dimension. Of course, I do not mean to suggest that all the learning that children do and that is relevant to their understanding of balance takes the form of explicit balance problems of the kind my network sees. Rather, my assumption that the experience on balance problems is dominated by problems in which there is no variability in weight is meant as a proxy for the more general assumption that children generally have more experience with weight than with distance as a factor in determining the relative heaviness of something.⁶ The specific assumptions about the sequence of learning experiences were as follows. The environment consisted of a list of training examples containing the full set of 625 possible

problems involving 25 combinations of possible weights (1-5 on the left crossed with 1-5 on the right) crossed with 25 combinations of possible distances (1-5 steps from the fulcrum on the left crossed with 1-5 steps from the fulcrum on the right). Two corpuses were set up. Problems in which the distance from the fulcrum was the same on both sides were listed five times each in one corpus, and 10 times each in the other corpus. Other problems were listed only once in each corpus.

Training and testing regime

Four simulation runs were carried out, two with each of the two corpora just described. In each run, training consisted of a series of epochs. In each epoch, 100 patterns were chosen randomly from the full list of patterns in the corpus. In each epoch, weight increments were accumulated over the 100 training trials and then added into the weights at the end of the epoch, according to the momentum method described in Rumelhart, Hinton, and Williams (1986, p. 330); parameters were $\eta = 0.075$, $\alpha = 0.9$.

After weight updating at the end of each epoch, the network was given a 24-item test, containing four problems of each of the six types described above, taken from an experiment of Siegler's. (A few of the examples had to be modified since Siegler's experiment had used up to six pegs.)

A comment on the simulation model

The model described above obviously simplifies the task that the learner faces and structures it for him to some degree. In particular, it embodies two principal assumptions which are crucial to the successful simulations we will consider below:

1. *Environment assumption.* The model assumes that the environment is biased, so that one dimension—in this case weight—is more frequently available as a basis for predicting outcome than the other.
2. *Architecture assumption.* The model assumes that the weight and distance dimensions are analysed separately, before information about the two dimensions is combined.

Both these assumptions are crucial to the success of the model. In an unbiased environment, both cues would be learned equally rapidly. Effects of combining the cues from the start, as prescribed by the architecture assumption, are more complex, but suffice it to say for now that the apparent stagelike character of performance is much less clear unless this assumption is adopted.

An important topic for further research will be to examine what variants of these assumptions might still allow the model to be successful. For example, regarding the environment, differences in salience (i.e. strength of input

activations) and structuredness of the dimensions might also produce similar results.

The issue of structuredness of the dimensions is a key point that needs to be considered as it relates to the present simulation. For both dimensions, the input representations encode different weights and distances from the fulcrum using distinct units. This means that different values are distinguishable by the model, but they are not structured for it; for example, the input itself provides no indication that a distance or weight of 3 is between 2 and 4. The network must learn to represent the weights and distances in structured ways in order to solve the balance problem. Below we will see that it does so.

Results

In general, performance of the model conformed to one of the four rules described by Siegler. Over the four runs, the model fit the criteria of one of Siegler's four rules on 85 per cent of the occasions, not counting an initial, pre-Rule 1 period discussed below (in Siegler, 1981, the conformity figure is about 90 per cent). Of course, the model was not consulting these rules or following the step-by-step procedures indicated in them; rather its behaviour was simply scorable by Siegler's criteria as consistent with the succession of rules. Excluding the initial period, failures to fit the rules were of three types: (1) cases in which a rule fit except for a position bias that gave difficulty on balance problems; (2) cases in which performance was borderline between Rules 1 and 2; and (3) combinations of these two problems. (Siegler (personal communication) does find some borderline cases between Rule 1 and Rule 2, but the position bias cases are not typical of children's performance.)

Overall development trends

Epoch by epoch performance in each of the four runs is shown in Figs 2.8 and 2.9. One generally observes the expected developmental progression. Each simulation run is slightly different, due to differences in the random starting weights and the sequence of actual training experiences, but there are clear common trends. Over the first 10 epochs or so, the output of the model was close to 0.5 on all test patterns; by our scoring criteria, all these outputs count as 'balance' responses, but of course they really represent a stage in which neither weight nor distance governs performance. The next few epochs represent a transition to Rule 1, in that in this phase the model is showing some tendency to activate the output unit on the side with the greater weight, but this tendency is variable across patterns and the discrepancy between the activations of the output units is not reliably greater than 0.33 when the weights differ.

After this brief transition, performance of the model had generally reached the point where it was responding consistently to the weight cue while

