

# Mind Design II

Philosophy  
Psychology  
Artificial Intelligence

Revised and enlarged edition

edited by  
John Haugeland

A Bradford Book  
The MIT Press  
Cambridge, Massachusetts  
London, England

*for Barbara and John III*

Second printing, 1997

© 1997 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

Book design and typesetting by John Haugeland. Body text set in Adobe Garamond 11.5 on 13; titles set in Zapf Humanist 601 BT. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Mind design II / edited by John Haugeland. — 2nd ed., rev. and enlarged.

p. cm.

"A Bradford book."

Includes bibliographical references.

ISBN 0-262-08259-4 (hc : alk. paper). — ISBN 0-262-58153-1 (pb : alk. paper)

1. Artificial intelligence. 2. Cognitive psychology.

I. Haugeland, John, 1945–

Q335.5.M492 1997

006.3—dc21

96-45188

CIP

not machines. Whatever else intentionality is, it is a biological phenomenon, and it is likely to be as causally dependent on the specific biochemistry of its origins as are lactation, photosynthesis, or any biological phenomena. No one would suppose that we could produce milk and sugar by running a computer simulation of the formal sequences in lactation and photosynthesis; but where the mind is concerned, many people are willing to believe in such a miracle, because of a deep and abiding dualism: the mind, they suppose, is a matter of formal processes and is independent of specific material causes in a way that milk and sugar are not.

In defense of this dualism, the hope is often expressed that the brain is a digital computer. (Early computers, by the way, were often called "electronic brains".) But that is no help. Of course the brain is a digital computer. Since everything is a digital computer, brains are too. The point is that the brain's causal capacity to produce intentionality cannot consist in its instantiating a computer program, since for any program you like it is possible for something to instantiate that program and still not have any mental states. Whatever it is that the brain does to produce intentionality, it cannot consist in instantiating a program, since no program by itself is sufficient for intentionality.

## Notes

1. I am not saying, of course, that Schank himself is committed to these claims.
2. Also, "understanding" implies both the possession of mental (intentional) states and the truth (validity, success) of these states. For the purposes of this discussion, we are concerned only with the possession of the states.
3. Intentionality is by definition that feature of certain mental states by which they are directed at or are about objects and states of affairs in the world. Thus, beliefs, desires, and intentions are intentional states; undirected forms of anxiety and depression are not. (For further discussion, see Searle 1979).

## The Architecture of Mind: A Connectionist Approach



David E. Rumelhart

1989

Cognitive science has a long-standing and important relationship to the computer. The computer has provided a tool whereby we have been able to express our theories of mental activity; it has been a valuable source of metaphors through which we have come to understand and appreciate how mental activities might arise out of the operations of simple-component processing elements.

I recall vividly a class I taught some fifteen years ago in which I outlined the then-current view of the cognitive system. A particularly skeptical student challenged my account, with its reliance on concepts drawn from computer science and artificial intelligence, with the question of whether I thought my theories would be different if it had happened that our computers were parallel instead of serial. My response, as I recall, was to concede that our theories might very well be different, but to argue that that wasn't a bad thing. I pointed out that the inspiration for our theories and our understanding of abstract phenomena always is based on our experience with the technology of the time. I pointed out that Aristotle had a wax tablet theory of memory, that Leibniz saw the universe as clockworks, that Freud used a hydraulic model of libido flowing through the system, and that the telephone-switchboard model of intelligence had played an important role as well. The theories posited by those of previous generations had, I suggested, been useful in spite of the fact that they were based on the metaphors of their time. Therefore, I argued, it was natural that in our generation—the generation of the serial computer—we should draw our insights from analogies with the most advanced technological developments of our time. I don't now remember whether my response satisfied the student, but I have no doubt that we in cognitive science have gained much of value through our use of concepts drawn from our experience with the computer.

In addition to its value as a source of metaphors, the computer differs from earlier technologies in another remarkable way. The computer can be made to *simulate* systems whose operations are very different from the computers on which these simulations run. In this way we can use the computer to simulate systems with which we *wish* to have experience and thereby provide a source of experience that can be drawn upon in giving us new metaphors and new insights into how mental operations might be accomplished. It is this use of the computer that the connectionists have employed. The architecture that we are exploring is not one based on the von Neumann architecture of our current generation of computers but rather an architecture based on considerations of how brains themselves might function. Our strategy has thus become one of offering a general and abstract model of the computational architecture of brains, to develop algorithms and procedures well suited to this architecture, to simulate these procedures and architecture on a computer, and to explore them as hypotheses about the nature of the human information-processing system. We say that such models are *neurally inspired*, and we call computation on such a system *brain-style computation*. Our goal in short is to replace the computer metaphor with the brain metaphor.

## 1 Why brain-style computation?

Why should a brain-style computer be an especially interesting source of inspiration? Implicit in the adoption of the computer metaphor is an assumption about the appropriate level of explanation in cognitive science. The basic assumption is that we should seek explanation at the *program* or *functional* level rather than the implementation level. Thus, it is often pointed out that we can learn very little about what kind of program a particular computer may be running by looking at the electronics. In fact we don't care much about the details of the computer at all; all we care about is the particular program it is running. If we know the program, we know how the system will behave in any situation. It doesn't matter whether we use vacuum tubes or transistors, whether we use an IBM or an Apple, the essential characteristics are the same. This is a very misleading analogy. It is true for computers because they are all essentially the same. Whether we make them out of vacuum tubes or transistors, and whether we use an IBM or an Apple computer, we are using computers of the same general design. But, when we look at an essentially different architecture, we see that the architecture makes

a good deal of difference. It is the architecture that determines which kinds of algorithms are most easily carried out on the machine in question. It is the architecture of the machine that determines the essential nature of the program itself. It is thus reasonable that we should begin by asking what we know about the architecture of the brain and how it might shape the algorithms underlying biological intelligence and human mental life.

The basic strategy of the connectionist approach is to take as its fundamental processing unit something close to an abstract neuron. We imagine that computation is carried out through simple interactions among such processing units. Essentially the idea is that these processing elements communicate by sending numbers along the lines that connect the processing elements. This identification already provides some interesting constraints on the kinds of algorithms that might underlie human intelligence.

The operations in our models then can best be characterized as "neurally-inspired". How does the replacement of the computer metaphor with the brain metaphor as model of mind affect our thinking? This change in orientation leads us to a number of considerations that further inform and constrain our model-building efforts. Perhaps the most crucial of these is time. Neurons are remarkably slow relative to components in modern computers. Neurons operate in the time scale of milliseconds, whereas computer components operate in the time scale of nanoseconds—a factor of  $10^6$  faster. This means that human processes that take on the order of a second or less can involve only a hundred or so time steps. Because most of the processes we have studied—perception, memory retrieval, speech processing, sentence comprehension, and the like—take about a second or so, it makes sense to impose what Feldman (1985a) calls the "100-step-program" constraint. That is, we seek explanations for these mental phenomena that do not require more than about a hundred elementary sequential operations. Given that the processes we seek to characterize are often quite complex and may involve consideration of large numbers of simultaneous constraints, our algorithms *must* involve considerable parallelism. Thus although a serial computer could be created out of the kinds of components represented by our units, such an implementation would surely violate the 100-step-program constraint for any but the simplest processes. Some might argue that, although parallelism is obviously present in much of human information processing, this fact alone need not greatly modify our world view. This is unlikely. The

speed of components is a critical design constraint. Although the brain has *slow* components, it has *very many* of them. The human brain contains billions of such processing elements. Rather than organize computation with many, many serial steps, as we do with systems whose steps are very fast, the brain must deploy many, many processing elements cooperatively and in parallel to carry out its activities. These design characteristics, among others, lead, I believe, to a general organization of computing that is fundamentally different from what we are used to.

A further consideration differentiates our models from those inspired by the computer metaphor—that is, the constraint that all the knowledge is *in the connections*. From conventional programmable computers we are used to thinking of knowledge as being stored in the states of certain units in the system. In our systems we assume that only very short-term storage can occur in the states of units; long-term storage takes place in the connections among units. Indeed it is the connections—or perhaps the rules for forming them through experience—that primarily differentiate one model from another. This is a profound difference between our approach and other more conventional approaches, for it means that almost all knowledge is *implicit* in the structure of the device that carries out the task, rather than *explicit* in the states of units themselves. Knowledge is not directly accessible to interpretation by some separate processor, but it is built into the processor itself and directly determines the course of processing. It is acquired through tuning of connections, as they are used in processing, rather than formulated and stored as declarative facts.

These and other neurally inspired classes of working assumptions have been one important source of assumptions underlying the connectionist program of research. These have not been the only considerations. A second class of constraints arises from our beliefs about the nature of human information processing considered at a more abstract, computational level of analysis. We see the kinds of phenomena we have been studying as products of a kind of constraint-satisfaction procedure in which a very large number of constraints act simultaneously to produce the behavior. Thus we see most behavior not as the product of a single, separate component of the cognitive system but as the product of a large set of interacting components, each mutually constraining the others and contributing in its own way to the globally observable behavior of the system. It is very difficult to use serial algorithms to implement such a conception but very natural to use highly

parallel ones. These problems can often be characterized as *best-match* or *optimization* problems. As Minsky and Papert (1969) have pointed out, it is very difficult to solve best-match problems serially. This is precisely the kind of problem, however, that is readily implemented using highly parallel algorithms of the kind we have been studying.

The use of brain-style computational systems, then, offers not only a hope that we can characterize how brains actually carry out certain information-processing tasks but also solutions to computational problems that seem difficult to solve in more traditional computational frameworks. It is here where the ultimate value of connectionist systems must be evaluated.

In this chapter, I begin with a somewhat more formal sketch of the computational framework of connectionist models. I then follow with a general discussion of the kinds of computational problems that connectionist models seem best suited for. Finally, I briefly review the state of the art in connectionist modeling.

## 1.1 The connectionist framework

There are seven major components of any connectionist system:

- a *set of processing units*;
- a *state of activation* defined over the processing units;
- an *output function* for each unit that maps its state of activation into an output;
- a *pattern of connectivity* among units;
- an *activation rule* for combining the inputs impinging on a unit with its current state to produce a new level of activation for the unit;
- a *learning rule* whereby patterns of connectivity are modified by experience; and
- an *environment* within which the system must operate.

Figure 8.1 illustrates the basic aspects of these systems. There is a set of processing units, generally indicated by circles in my diagrams; at each point in time each unit  $u_i$  has an activation value, denoted in the diagram as  $a_i(t)$ ; this activation value is passed through a function  $f_i$  to produce an output value  $o_i(t)$ . This output value can be seen as passing through a set of unidirectional connections (indicated by lines or arrows in the diagrams) to other units in the system. There is associated with each connection a real number, usually called the *weight* or

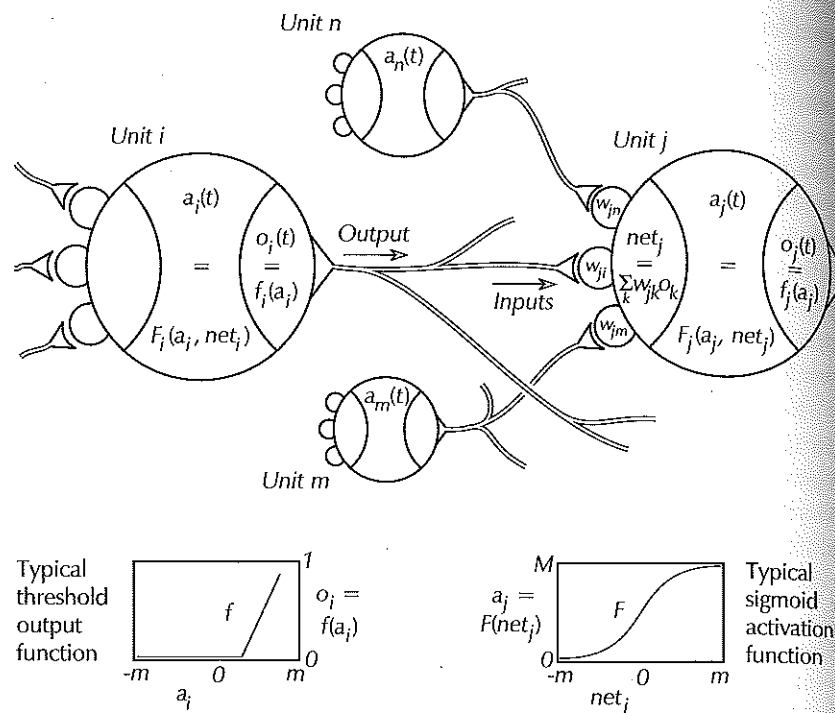


Figure 8.1: The basic parts of a parallel distributed processing system.

*strength* of the connection, designated  $w_{ij}$  (to unit  $i$ , from unit  $j$ ), which determines how strongly the former is affected by the latter. All of the inputs must then be combined; and the combined inputs to a unit (usually designated the *net input* to that unit), along with its current activation value, determine its new activation value via a function  $F$ . These systems are viewed as being plastic in the sense that the pattern of interconnections is not fixed for all time; rather the weights can undergo modification as a function of experience. In this way the system can evolve. What a unit represents can change with experience, and the system can come to perform in substantially different ways.

**THE SET OF PROCESSING UNITS.** Any connectionist system begins with a set of processing units. Specifying the set of processing units and what they represent is typically the first stage of specifying a connectionist model. In some systems these units may represent particular conceptual objects such as features, letters, words, or concepts; in others they are simply abstract elements over which meaningful patterns can be defined. When we speak of a distributed representation, we

mean one in which the units represent small, featurelike entities we call *microfeatures*. In this case it is the pattern as a whole that is the meaningful level of analysis. This should be contrasted to a *one-unit-one-concept* or *localist* representational system, in which single units represent entire concepts or other large meaningful entities.

All of the processing of a connectionist system is carried out by these units. There is no executive or other overseer. There are only relatively simple units, each doing its own relatively simple job. A unit's job is simply to receive input from its neighbors and, as a function of the inputs it receives, to compute an output value, which it sends to its neighbors. The system is inherently parallel in that many units can carry out their computations at the same time.

Within any system we are modeling, it is useful to characterize three types of units: *input*, *output*, and *hidden* units. Input units receive inputs from sources external to the system under study. These inputs may be either sensory inputs or inputs from other parts of the processing system in which the model is embedded. The output units send signals out of the system. They may either directly affect motoric systems or simply influence other systems external to the ones we are modeling. The hidden units are those whose only inputs and outputs are within the system we are modeling. They are not "visible" to outside systems.

**THE STATE OF ACTIVATION.** In addition to the set of units we need a representation of the state of the system at time  $t$ . This is primarily specified by a vector  $\mathbf{a}(t)$ , representing the pattern of activation over the set of processing units. Each element of the vector stands for the activation of one of the units. It is the *pattern* of activation over the whole set of units that captures what the system is representing at any time. It is useful to see processing in the system as the evolution, through time, of a pattern of activity over the set of units.

Different models make different assumptions about the activation values a unit is allowed to take on. Activation values may be continuous or discrete. If they are continuous, they may be unbounded or bounded. If they are discrete, they may take binary values or any of a small set of values. Thus in some models units are continuous and may take on any real number as an activation value. In other cases they may take on any real value between some minimum and maximum such as, for example, the interval  $[0, 1]$ . When activation values are restricted to discrete values, they most often are binary—such as the values 0 and 1,

where 1 is usually taken to mean that the unit is active and 0 is taken to mean that it is inactive.

**THE OUTPUT FUNCTION.** Units interact by transmitting signals to their neighbors. The strengths of their signals, and therefore the degrees to which they affect their neighbors, are determined by their levels of activation. Associated with each unit  $u_i$  is an output function  $f_i(a_i(t))$ , which maps the current state of activation to an output signal  $o_i(t)$ . In some of our models, the output level is exactly equal to the activation level of the unit. In this case,  $f$  is the identity function  $f(x) = x$ . Sometimes  $f$  is some sort of threshold function, so that a unit has no effect on another unit unless its activation exceeds a certain value. Sometimes the function  $f$  is assumed to be a stochastic function in which the output of the unit depends probabilistically on its activation level.

**THE PATTERN OF CONNECTIVITY.** Units are connected to one another. It is this pattern of connectivity that constitutes what the system knows and determines how it will respond to any arbitrary input. Specifying the processing system and the knowledge encoded therein is, in a connectionist model, a matter of specifying this pattern of connectivity among the processing units.

In many cases we assume that each unit provides an additive contribution to the input of the units to which it is connected. In such cases the total input to any unit is simply the weighted sum of the separate inputs from each of the units connected to it. That is, the inputs from all of the incoming units are simply multiplied by their respective connection weights and summed to get the overall input to that unit. In this case the total pattern of connectivity can be represented by merely specifying the weights for each of the connections in the system. A positive weight represents an excitatory input, and a negative weight represents an inhibitory input. It is often convenient to represent such a pattern of connectivity by a weight matrix  $\mathbf{W}$  in which the entry  $w_{ij}$  represents the strength and sense of the connection to unit  $u_i$  from unit  $u_j$ . The weight  $w_{ij}$  is a positive number if unit  $u_j$  excites unit  $u_i$ ; it is a negative number if unit  $u_j$  inhibits unit  $u_i$ ; and it is 0 if unit  $u_j$  has no direct connection to unit  $u_i$ . The absolute value of  $w_{ij}$  specifies the *strength of the connection*.

The pattern of connectivity is very important. It is this pattern that determines what each unit represents. One important issue that may determine both how much information can be stored and how much

serial processing the network must perform is the *fan-in* and *fan-out* of a unit. The fan-in is the number of elements that either excite or inhibit a given unit. The fan-out is the number of units affected directly by a unit. It is useful to note that in brains these numbers are relatively large. Fan-in and fan-out range as high as 100,000 in some parts of the brain. It seems likely that this large fan-in and fan-out allows for a kind of operation that is less like a fixed circuit and more statistical in character.

**THE ACTIVATION RULE.** We also need a rule whereby the inputs impinging on a particular unit are combined with one another and with the current state of the unit to produce a new state of activation. We need a function  $F_i$ , which takes  $a_i(t)$  and the net inputs,  $\text{net}_i = \sum_j w_{ij} o_j(t)$ , and produces a new state of activation. In the simplest cases, when  $F_i$  is the identity function and depends only on the inputs, we can write  $a_i(t+1) = \text{net}_i(t)$  — or, in vector notation for the whole network at once,  $\mathbf{a}(t+1) = \mathbf{net}(t) = \mathbf{W}\mathbf{o}(t)$ . Sometimes  $F$  is a threshold function so that the net input must exceed some value before contributing to the new state of activation. Often the new state of activation depends on the old one as well as the current input. The function  $F$  itself is what we call the activation rule. Usually the function is assumed to be deterministic. Thus, for example, if a threshold is involved, it may be that  $a_i(t) = 1$  if the total input exceeds some threshold value, and equals 0 otherwise. Other times it is assumed that  $F$  is stochastic. Sometimes activations are assumed to decay slowly with time so that even with no external input the activation of a unit will simply decay and not go directly to zero. Whenever  $a_i(t)$  is assumed to take on continuous values, it is common to assume that  $F$  is a kind of sigmoid (that is, S-shaped) function. In this case an individual unit can *saturate* and reach a minimum or maximum value of activation.

**THE LEARNING RULE: CHANGES AS A FUNCTION OF EXPERIENCE.** Changing the processing or knowledge structure in a connectionist system involves modifying the patterns of interconnectivity. In principle this can involve three kinds of modification:

- (1) development of new connections;
- (2) loss of existing connections;
- (3) modification of the strengths of connections that already exist.

Very little work has been done on (1) and (2). To a first order of approximation, however, (1) and (2) can be considered a special case of

(3). Whenever we change the strength of connection away from zero to some positive or negative value, it has the same effect as growing a new connection. Whenever we change the strength of a connection to zero, that has the same effect as losing an existing connection. Thus we have concentrated on rules whereby *strengths* of connections are modified through experience.

Virtually all learning rules for models of this type can be considered variants of the *Hebbian* learning rule, suggested by Hebb in his classic book *Organization of Behavior* (1949). Hebb's basic idea is this: if a unit  $u_i$  receives an input from another unit  $u_j$  at a time when both units are highly active, then the weight  $w_{ij}$  to  $u_i$  from  $u_j$  should be *strengthened*. This idea has been extended and modified so that it can be stated more generally as

$$\delta w_{ij} = g(a_i(t), \tau_i(t)) \cdot h(o_j(t), w_{ij})$$

or, suppressing the time variables for easier readability, as

$$\delta w_{ij} = g(a_i, \tau_i) \cdot h(o_j, w_{ij})$$

where  $\tau_i$  is a kind of *teaching* input to  $u_i$ . Simply stated, this equation says that the *change* in the connection to  $u_i$  from  $u_j$  is given by the product of a function  $g(\dots)$  of the activation of  $u_i$  and its teaching input  $\tau_i$  and another function  $h(\dots)$  of the output value of  $u_j$  and the current connection strength  $w_{ij}$ . In the simplest versions of Hebbian learning, there is no teacher and the functions  $g$  and  $h$  are simply proportional to their first arguments. Thus we have

$$\delta w_{ij} = \epsilon a_i o_j$$

where  $\epsilon$  is the constant of proportionality representing the learning rate. Another common variation is a rule in which

$$h(o_j, w_{ij}) = o_j$$

(as in the simplest case) but

$$g(a_i, \tau_i) = \epsilon \cdot (\tau_i - a_i)$$

This is often called the *Widrow-Hoff* rule, because it was originally formulated by Widrow and Hoff (1960), or the *delta* rule, because the amount of learning is proportional to the *difference* (or delta) between the actual activation achieved and the target activation provided by a teacher. In this case we have

$$\delta w_{ij} = \epsilon \cdot (\tau_i - a_i) \cdot o_j$$

This is a generalization of the *perceptron* learning rule for which the famous *perceptron convergence theorem* has been proved. Still another variation has

$$\delta w_{ij} = \epsilon a_i \cdot (o_i - w_{ij})$$

This is a rule employed by Grossberg (1976) and others in the study of *competitive learning*. In this case usually only the units with the strongest activation values are allowed to learn.

**THE ENVIRONMENT.** It is crucial in the development of any model to have a clear representation of the environment in which this model is to exist. For connectionist models, we represent the environment as a time-varying stochastic function over a space of possible input patterns. That is, for each possible input pattern, we imagine that there is some probability that, at any given time, that pattern is impinging on the input units. This probability function may in general depend on the history of inputs to the system as well as outputs of the system. In practice most connectionist models involve a much simpler characterization of the environment. Typically, the environment is characterized by a stable probability distribution over the set of possible input patterns, independent of past inputs and past responses of the system. In this case we can imagine listing the set of possible inputs to the system and numbering them from 1 to  $M$ . The environment is then characterized by a set of probabilities  $p_i$  for  $i = 1, \dots, M$ . Because each input pattern can be considered a vector, it is sometimes useful to characterize those patterns with nonzero probabilities as constituting *orthogonal* or *linearly independent* sets of vectors.

To summarize, the connectionist framework consists not only of a formal language but also a perspective on our models. Other qualitative and quantitative considerations arising from our understanding of brain processing and of human behavior combine with the formal system to form what might be viewed as an aesthetic for our model-building enterprises.

## 1.2 Computational features of connectionist models

In addition to the fact that connectionist systems are capable of exploiting parallelism in computation and mimicking brain-style computation, such systems are important because they provide good solutions to a number of very difficult computational problems that seem to arise often in models of cognition. In particular they typically:

- are good at solving constraint-satisfaction problems;
- are efficient mechanisms for best-match search, pattern recognition, and content-addressable memory;
- automatically implement similarity-based generalization;
- offer simple, general mechanisms for adaptive learning; and
- exhibit graceful degradation with damage or information overload.

**CONSTRAINT SATISFACTION PROBLEMS.** Many cognitive-science problems are usefully conceptualized as problems in which a solution is given through the satisfaction of a very large number of mutually interacting constraints. The challenge is to devise a computational system that is capable of efficiently solving such problems. Connectionist networks are ideal for implementing constraint-satisfaction systems; indeed, the trick for getting connectionist networks to solve difficult problems is often to cast the problems as constraint-satisfaction problems. In this case, we conceptualize the connectionist network as a *constraint network* in which each unit represents a hypothesis of some sort (for example, that a certain semantic feature, visual feature, or acoustic feature is present in the input), and each connection represents a constraint among the hypotheses.

Thus, for such a network, if feature B is expected to be present whenever feature A is present, there should be a positive connection from the unit corresponding to the hypothesis that A is present to the unit representing the hypothesis that B is present. Contrariwise, if there is a constraint that whenever A is present B is expected *not* to be present, there should be a negative connection from A to B. If the constraints are weak, the weights should be small; if the constraints are strong, then the weights should be large. Similarly, the inputs to such networks can also be thought of as constraints. A positive input to a particular unit means that there is evidence from the outside that the relevant feature is present. A negative input means that there is evidence from the outside that the feature is not present. The stronger the input, the greater the evidence. If a network of this kind is allowed to run, it will eventually *settle* into an optimal state in which as many as possible of the constraints are satisfied, with priority given to the strongest constraints. (Actually, the system will find a *locally* best solution to the constraint-satisfaction problem. *Global* optima are more difficult to find.) The procedure whereby such a system *settles* into such a state

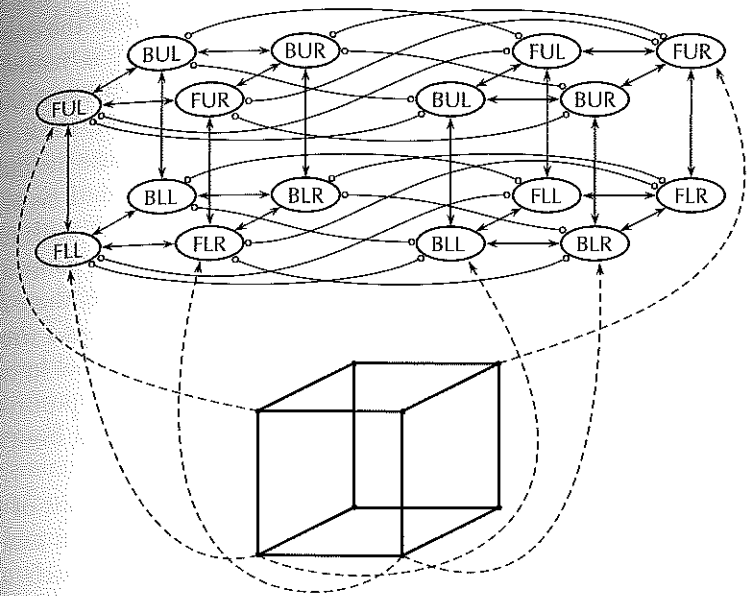


Figure 8.2: A simple network representing some constraints involved in perceiving a Necker cube. The ovals are the units in the network; connections with arrow-heads are positive (excitatory), while those with circle-heads are negative (inhibitory); the dotted lines represent input stimuli from the perceived cube.

is called *relaxation*. We speak of the system *relaxing* to a solution. Thus, many connectionist models are constraint-satisfaction models that settle on locally optimal solutions through a process of relaxation.

Figure 8.2 shows an example of a simple 16-unit constraint network. Each unit in the network represents a hypothesis concerning a vertex in a line drawing of a Necker cube. The network consists of two interconnected subnetworks—one corresponding to each of the two global interpretations of the Necker cube. Each unit in each subnetwork is assumed to receive input from the region of the input figure—the cube—corresponding to its location in the network. Each unit in figure 8.2 is labeled with a three-letter sequence indicating whether its vertex is hypothesized to be front or back (F or B), upper or lower (U or L), and right or left (R or L). Thus, for example, the lower-left unit of each subnetwork is assumed to receive input from the lower-left vertex of the input figure. The unit in the left network represents the hypothesis that it is receiving input from a lower-left vertex in the

front surface of the cube (and is thus labeled FLL), whereas the one in the right subnetwork represents the hypothesis that it is receiving input from a lower-left vertex in the back surface (BLL). Because there is a constraint that each vertex has a single interpretation, these two units are connected by a strong negative connection. Because the interpretation of any given vertex is constrained by the interpretations of its neighbors, each unit in a subnetwork is connected positively with each of its neighbors within the network. Finally, since there is a constraint that there can be only one vertex of each kind (for example, there can be only one lower-left vertex in the front plane, FLL), there is a strong negative connection between units representing the same label in each subnetwork. Thus each unit has three neighbors connected positively, two competitors connected negatively, and one positive input from the stimulus.

For purposes of this example, we assume that the strengths of the connections have been arranged so that two negative inputs exactly balance three positive inputs. Further it is assumed that each unit receives an excitatory input from the ambiguous stimulus pattern and that each of these excitatory influences is relatively small. Thus, if all three of a unit's neighbors are on and both of its competitors are on, these effects would entirely cancel out one another; and if there were a small input from the outside, the unit would have a tendency to come on. On the other hand, if fewer than three of its neighbors were on and both of its competitors were on, the unit would have a tendency to turn off, even with an excitatory input from the stimulus pattern.

In the preceding paragraphs, I focused on the individual units of the networks. It is often useful, however, to focus not on the units but on entire *states* of the network. In the case of binary (on-off or 0-1) units, there would be a total of  $2^{16}$  possible states in which a network of this size could reside—since, in principle, each of the 16 units could have either value 0 or 1. In the case of continuous units, in which each unit can take on any value between 0 and 1, the system could in principle take on any of an infinite number of states. Yet because of the constraints built into the network, there are only a relatively few of those states into which the system will ever actually settle.

To see this, consider the case in which the units are updated asynchronously, one at a time. During each time slice, one of the units is chosen to update. If its net input exceeds 0, its value will be pushed toward 1; otherwise its value will be pushed toward 0. Imagine that the system starts with all units off. A unit is then chosen at random to be

updated. Because it is receiving slight positive input from the stimulus and no other inputs, it will be given a positive activation value. Then another unit is chosen to update. Unless it is in direct competition with the first unit, it too will be turned on. Eventually a coalition of neighboring units will be turned on. These units will tend to turn on more of their neighbors in the same subnetwork and turn off their competitors in the other subnetwork. The system will (almost always) end up in a situation in which all of the units in one subnetwork are fully activated and none of the units in the other subnetwork is activated. That is, the system will end up interpreting the Necker cube as either facing left or facing right. Whenever the system gets into a state and stays there, the state is called a *stable state* or a *fixed point* of the network. The constraints implicit in the pattern of connections among the units determine the set of possible stable states of the system and therefore the set of possible interpretations of the inputs.

Hopfield (1982) has shown that it is possible to give a general account of the behavior of systems like this (with symmetric weights and asynchronous updates). In particular, he has shown that such systems can be conceptualized as minimizing a global measure, which he calls the *energy* of the system, through a method of *gradient descent* or, equivalently, maximizing the constraints satisfied through a method of *hill climbing*. More specifically, the system operates so as to move always from a state that satisfies fewer constraints to one that satisfies more, where the measure of constraint satisfaction is given by

$$G(t) = \sum_i \sum_j w_{ij} a_i(t) a_j(t) + \sum_i \text{input}_i(t) a_i(t)$$

Essentially the equation says that the overall goodness of fit is given by the sum of the degrees to which each pair of units contributes to the goodness plus the degree to which the units satisfy the input constraints. The contribution of a pair of units is given by the product of their activation values and the weight connecting them. Thus, if the weight is positive, each unit wants to be as active as possible—that is, the activation values for those two units should be pushed toward 1. If the weight is negative, then at least one of the units should be 0 to maximize the pairwise goodness. Similarly if the input constraint for a given unit is positive, then its contribution to the total goodness of fit is maximized by bringing the activation of that unit toward its maximal value. If it is negative, the activation value should be decreased toward 0. Of course the constraints will generally not be

totally consistent. Sometimes a given unit may have to be turned on to increase the function in some ways yet decrease it in other ways. The point is that it is the sum of all of these individual contributions that the system seeks to maximize. Thus, for every state of the system—every possible pattern of activation over the units—the pattern of inputs and the connectivity matrix  $\mathbf{W}$  determine a value of the goodness-of-fit function. The system processes its input by moving upward from state to adjacent state until it reaches a state of maximum goodness. When it reaches such a *stable state* or *fixed point*, it will stay in that state and it can be said to have “settled” on a solution to the constraint-satisfaction problem or, as in our present case, “settled into an interpretation” of the input.

It is important to see then that entirely *local* computational operations, in which each unit adjusts its activation up or down on the basis of its net input, serve to allow the network to converge toward states that maximize a *global* measure of goodness or degree of constraint satisfaction. Hopfield's main contribution to the present analysis was to point out this basic fact about the behavior of networks with symmetrical connections and asynchronous update of activations.

Finally, one of the most difficult problems in cognitive science is to build systems that can allow a large number of knowledge sources to interact usefully in the solution of a problem. Thus, in language processing we would want syntactic, phonological, semantic, and pragmatic knowledge sources all to interact in the construction of the meaning of an input. Reddy and his colleagues (1973) have had some success in the case of speech perception with the Hearsay system because they were working in the highly structured domain of language. Less structured domains have proved very difficult to organize. Connectionist models, conceived as constraint-satisfaction networks, are ideally suited for blending multiple-knowledge sources. Each knowledge type is simply another constraint, and the system will, in parallel, find those configurations of values that best satisfy all of the constraints from all of the knowledge sources. The uniformity of representation and the common currency of interaction (activation values) make connectionist systems especially powerful for this domain.

To summarize, there is a large subset of connectionist models that can be considered constraint-satisfaction models. These networks can be described as carrying out their information processing by climbing into states of maximal satisfaction of the constraints implicit in the network. A very useful consequence of this way of viewing networks is

that we can describe their behavior not only in terms of the behavior of individual units but also in terms of the properties of the network itself. A primary concept for understanding these network properties is the *goodness-of-fit landscape* over which the system moves. Once we have correctly described this landscape, we have described the operational properties of the system—it will process information by moving uphill toward goodness maxima. The particular maximum that the system will find is determined by where the system starts and by the distortions of the space induced by the input. One of the very important descriptors of a goodness landscape is the set of maxima that a system can find, the size of the region that feeds into each maximum, and the height of the maximum itself. The states themselves correspond to possible interpretations, the peaks in the space correspond to the best interpretations, the extent of the foothills or skirts surrounding a particular peak determines the likelihood of finding the peak, and the height of the peak corresponds to the degree to which the constraints of the network are actually met or alternatively to the goodness of the interpretation associated with the corresponding state.

**BEST-MATCH SEARCH, PATTERN RECOGNITION, AND CONTENT-ADDRESSABLE MEMORY.** These are all variants on the general best-match problem (compare Minsky and Papert 1969). Best-match problems are especially difficult for serial computational algorithms (they involve exhaustive search), but, as we have just indicated, connectionist systems can readily be used to find the interpretation that best matches a set of constraints.

They can similarly be used to find stored data that best match some target or probe. In this case, it is useful to imagine that the network consists of two classes of units. One class, the *visible* units, corresponds to the contents stored in the network, in the sense that each stored pattern is a possible pattern of activation of these units. The other units, the *hidden* units, correspond to shared structural properties of the stored patterns that play a role in storing and retrieving them. The patterns themselves are actually stored in the weights on the connections among all these units. If we think of each stored pattern as a collection of features, then each visible unit corresponds to the hypothesis that some particular feature is present in the relevant pattern, and each hidden unit corresponds to a hypothesis concerning a *configuration* of several features. The hypothesis to which a particular hidden unit corresponds is determined by the exact *learning rule* used to store the

input and by the characteristics of the ensemble of stored patterns. Retrieval in such a network amounts to setting the values of *some* of the visible units (the retrieval probe) and letting the system settle to the best interpretation of that input, while itself setting the values of the remaining visible units. This is a kind of pattern completion. The details are not too important here because a variety of learning rules lead to networks that all have the following important properties:

- When a previously stored (that is, familiar) pattern enters the memory system, it is amplified, and the system responds with a stronger version of the input pattern. This is a kind of recognition response.
- When an unfamiliar pattern enters the memory system, it is dampened, and the activity of the memory system is shut down. This is a kind of unfamiliarity response.
- When part of a familiar pattern is presented, the system responds by "filling in" the missing parts. This is a kind of recall paradigm in which the part constitutes the retrieval cue, and the filling in is a kind of memory-reconstruction process. This is a content-addressable memory system.
- When a pattern similar to a stored pattern is presented, the system responds by distorting the input pattern toward the stored pattern. This is a kind of assimilation response in which similar inputs are assimilated to similar stored events.
- Finally, if a number of similar patterns have been stored, the system will respond strongly to the central tendency of the stored patterns, even though the central tendency itself was never stored. Thus this sort of memory system automatically responds to prototypes even when no prototype has been seen.

These properties correspond very closely to the characteristics of human memory and, I believe, are exactly the kind of properties we want in any theory of memory.

**AUTOMATIC, SIMILARITY-BASED GENERALIZATION.** One of the major complaints against AI programs is their "fragility". The programs are usually very good at what they are programmed to do, but respond in unintelligent or odd ways when faced with novel situations. There seem to be at least two reasons for this fragility. In conventional symbol-processing systems similarity is represented only indirectly, and is therefore not available as a basis for generalizations; and most AI programs are not self-modifying and cannot adapt to their environments.

In our connectionist systems, on the other hand, similarities among patterns are directly represented along with the patterns themselves in the connection weights—in such a way that similar patterns have similar effects. Therefore, similarity-based generalization is an automatic property of connectionist models. It should be noted that the degree of similarity between patterns is roughly given by the inner product of the vectors representing the patterns. Thus the dimensions of generalization are given by the dimensions of the representational space. Often this will lead to the right generalizations. But, there are situations in which it leads to inappropriate generalizations. In such cases, we must allow the system to *learn* its appropriate representation. In the next section I describe how the appropriate representation can be learned so that the correct generalizations are automatically made.

**LEARNING.** A key advantage of connectionist systems is the fact that simple yet powerful learning procedures can be defined that allow the systems to adapt to their environments. It was work on the learning aspect of neurally inspired models that first led to an interest in them (compare Rosenblatt 1962), and it was the demonstration that those learning procedures could not work for complex networks that contributed to the loss of interest (compare Minsky and Papert 1969). Although the *perceptron convergence procedure* and its variants have been around for some time, they are limited to simple two-layer networks involving only input and output units. There were no hidden units in these cases and no internal representation. The coding provided by the external world had to suffice. Nevertheless, such networks have proved useful in a wide variety of applications. Perhaps their most important characteristic is that they map similar input patterns to similar output patterns. This is what allows them to make reasonable generalizations and perform reasonably on patterns that have never before been presented. The similarity of patterns in connectionist systems is determined by their overlap. This overlap, for two-layer networks, is determined entirely outside the learning system itself—by whatever produces the patterns.

The constraint that similar input patterns lead to similar outputs can lead to an inability of the system to learn certain mappings from input to output. Whenever the representation provided by the outside world is such that the similarity structure of the input and output patterns is very different, a network without internal representations (that is, a network without hidden units) will be unable to perform the

Input patterns		Output patterns
00	→	0
01	→	1
10	→	1
11	→	0

Table 8.1: XOR problem.

necessary mappings. A classic example of this case is the exclusive-or (XOR) problem illustrated in table 8.1. Here we see that those patterns that overlap least are supposed to generate identical output values. This problem and many others like it cannot be solved by networks that lack hidden units with which to create their own internal representations of the input patterns. It is interesting to note that if the input patterns contained a third input bit, taking the value 1 when and only when the other two were both 1 (as shown in table 8.2), a two-layer system would be able to solve the problem.

Minsky and Papert (1969) have provided a careful analysis of conditions under which such systems are capable of carrying out the required mappings. They show that in many interesting cases networks of this kind are incapable of solving the problems. On the other hand, as Minsky and Papert also point out, if there is a layer of simple perceptron-like hidden units, as shown in figure 8.3, with which the original input pattern can be augmented, there is always a recoding (that is, an internal representation) of the input patterns in the hidden units in which the similarity of the patterns among the hidden units can support any required mapping from the input to the output units. Thus if we have the right connections from the input units to a large enough set of hidden units, we can always find a representation that will perform any mapping from input to output through these hidden units. In the case of the XOR problem, the addition of a feature that detects the conjunction of the input units changes the similarity structure of the patterns sufficiently to allow the solution to be learned. As illustrated in figure 8.4, this can be done with a single hidden unit. The numbers on the arrows represent the strengths of the connections among the units. The numbers written in the circles represent the thresholds of the units. The value of +1.5 for the threshold of the hidden unit ensures that it will be turned on only when both input units

Input patterns		Output patterns
000	→	0
010	→	1
100	→	1
111	→	0

Table 8.2: XOR problem with redundant third bit.

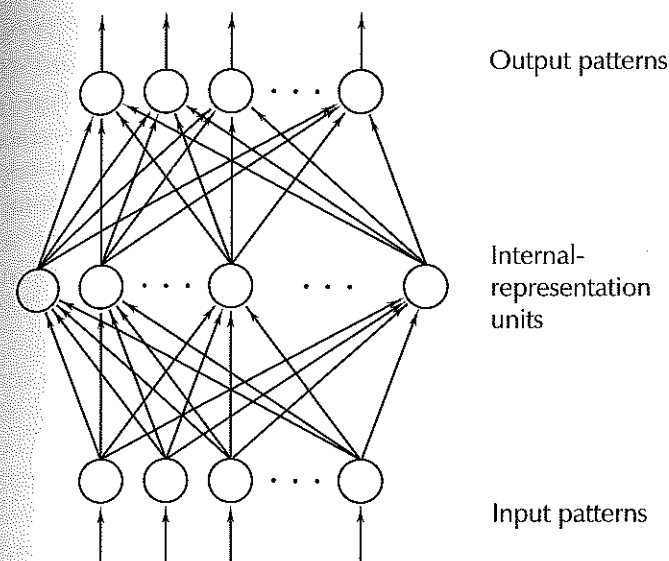


Figure 8.3: A multilayer network in which input patterns are recoded by internal representation units.

are on. The threshold of 0.5 for the output unit ensures that it will turn on only when it receives a net positive input greater than 0.5. The weight of -2 from the hidden unit to the output unit ensures that the output unit will not come on when both input units are on. Note that, from the point of view of the output unit, the hidden unit is treated as simply another input unit. It is as if the input patterns consisted of three rather than two units (essentially as in table 8.2).

The existence of networks such as this illustrates the potential power of hidden units and internal representations. The problem, as

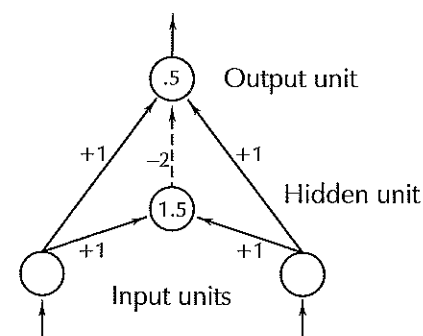


Figure 8.4: A simple XOR network with one hidden unit.

noted by Minsky and Papert, is that, whereas there is a very simple guaranteed learning rule for all problems that can be solved without hidden units—namely, the perceptron convergence procedure (or the variation reported originally by Widrow and Hoff 1960)—there has been no equally powerful rule for learning in multilayer networks.

It is clear that if we hope to use these connectionist networks for general computational purposes, we must have a learning scheme capable of learning its own internal representations. This is just what we (Rumelhart, Hinton, and Williams 1986) have done. We have developed a generalization of the perceptron learning procedure, called the *generalized delta rule*, which allows the system to learn to compute arbitrary functions. The constraints inherent in networks without self-modifying internal representations are no longer applicable. The basic learning procedure is a two-stage process. First, an input is applied to the network. Then, after the system has processed for some time, certain units of the network—usually the output units—are informed of the values they ought to have attained. If they have attained the desired values, the weights on their input connections are left unchanged. If they differ from their target values, then those weights are changed slightly, in such a way as to reduce the differences between the actual values attained and the target values.

Those *differences* between the actual and target values at the output units can be thought of as *error signals*. Similar error signals must be sent back in turn to those units that impinged on the output units. Each such unit receives an error signal that is equal to the sum of the errors in each of the output units to which it connects times the weight on the connection to that output unit. Then, based on those error signals, the weights on the *input* connections into those “second-layer” units can be modified, after which error signals can be passed back another layer. This process—called the *backpropagation of error*—continues until the error signals reach the input units or until they have been passed back a predetermined number of times. Then a new input pattern is presented and the process repeats. Although the procedure may sound difficult, it is actually quite simple and easy to implement within these nets. As shown in Rumelhart, Hinton, and Williams (1986), such a procedure will always change its weights in such a way as to reduce the overall difference between the actual output values and the desired output values. Moreover it can be shown that this system will work for any network whatsoever.

Minsky and Papert, in their pessimistic discussion of perceptrons, discuss *multilayer machines*. They state that

The perceptron has shown itself worthy of study despite (and even because of!) its severe limitations. It has many features that attract attention: its linearity; its intriguing learning theorem; its clear paradigmatic simplicity as a kind of parallel computation. There is no reason to suppose that any of these virtues carry over to the many-layered version. Nevertheless, we consider it to be an important research problem to elucidate (or reject) our intuitive judgment that the extension is sterile. Perhaps some powerful convergence theorem will be discovered, or some profound reason for the failure to produce an interesting “learning theorem” for the multilayered machine will be found. (1969, pp. 231–232)

Although our learning results do not *guarantee* that we can find a solution for all solvable problems, our analysis and simulation results have shown that, as a practical matter, this error-propagation scheme leads to solutions in virtually every case. In short, I believe that we have answered Minsky and Papert’s challenge and *have* found a learning result sufficiently powerful to demonstrate that their pessimism about learning in multilayer machines was misplaced.

One way to view the procedure I have been describing is as a parallel computer that, having been shown the appropriate input/output exemplars specifying some function, programs itself to compute that function in general. Parallel computers are notoriously difficult to program. Here we have a mechanism whereby we do not actually have to know how to write the program to get the system to do it.

**GRACEFUL DEGRADATION.** Finally, connectionist models are interesting candidates for cognitive-science models because of their property of graceful degradation in the face of damage and information overload. The ability of our networks to learn leads to the promise of computers that can literally learn their way around faulty components: because every unit participates in the storage of many patterns and because each pattern involves many different units, the loss of a few components will degrade the stored information, but will not destroy it. Similarly such memories should not be conceived as having a certain fixed capacity. Rather, there is simply more and more storage interference and blending of similar pieces of information as the memory is overloaded. This property of graceful degradation mimics the

human response in many ways and is one of the reasons we find these models of human information processing plausible.

## 2 The state of the art

Recent years have seen a virtual explosion of work in the connectionist area. This work has been singularly interdisciplinary, being carried out by psychologists, physicists, computer scientists, engineers, neuroscientists, and other cognitive scientists. A number of national and international conferences have been established and are being held each year. In such an environment it is difficult to keep up with the rapidly developing field. Nevertheless, a reading of recent papers indicates a few central themes in this activity. These themes include the study of learning and generalization (especially the use of the backpropagation learning procedure), applications to neuroscience, mathematical properties of networks—both of the learning algorithms and of connectionist style computation itself in comparison to more conventional computational paradigms—and finally the development of an implementational base for physical realizations of connectionist computational devices, especially in the areas of optics and analog VLSI.

Although there are many other interesting and important developments, I conclude with a brief summary of the work with which I have been most involved over the past several years, namely, the study of learning and generalization within multilayer networks. Even this summary is necessarily selective, but it should give a sampling of much of the current work in the area.

The backpropagation learning procedure has become possibly the single most popular method for training networks. The procedure has been used to train networks on problem domains including character recognition, speech recognition, sonar detection, mapping from spelling to sound, motor control, analysis of molecular structure, diagnosis of eye diseases, prediction of chaotic functions, playing backgammon, the parsing of simple sentences, and many, many more areas. Perhaps the major point of these examples is the enormous range of problems to which the backpropagation learning procedure can usefully be applied. In spite of the rather impressive breadth of topics and the success of some of these applications, there are a number of serious open problems. The theoretical issues of primary concern fall into three main areas. (1) The architecture problem: are there useful architectures beyond the standard three-layer network that are appropriate for

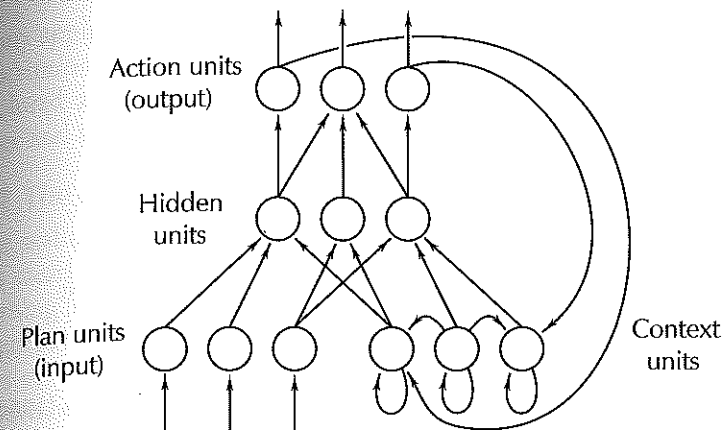


Figure 8.5: A recurrent network of the type developed by Jordan (1986) for learning to perform sequential operations.

certain areas of application? (2) The scaling problem: how can we cut down on the substantial training time that seems to be involved for more difficult and interesting problems? (3) The generalization problem: how can we be certain that a network trained on a subset of the example set will generalize correctly to the entire set of examples?

### 2.1 Some architecture

Although most applications have involved the simple three-layer backpropagation network with one input layer, one hidden layer, and one output layer of units, there have been a large number of interesting architectures proposed—each for the solution of some particular problem of interest. There are, for example, a number of “special” architectures that have been proposed for the modeling of such sequential phenomena as motor control. Perhaps the most important of these is the one proposed by Mike Jordan (1986) for producing sequences of phonemes. The basic structure of the network is illustrated in figure 8.5. It consists of four groups of units. *Plan units*, which tell the network which sequence it is producing, are fixed at the start of a sequence and are not changed. *Context units*, which keep track of where the system is in the sequence, receive input from the output units of the systems and from themselves, constituting a memory for the sequence produced thus far. *Hidden units* combine the information from the plan units with that from the context units to determine which output is to be produced next. *Output units* produce the desired

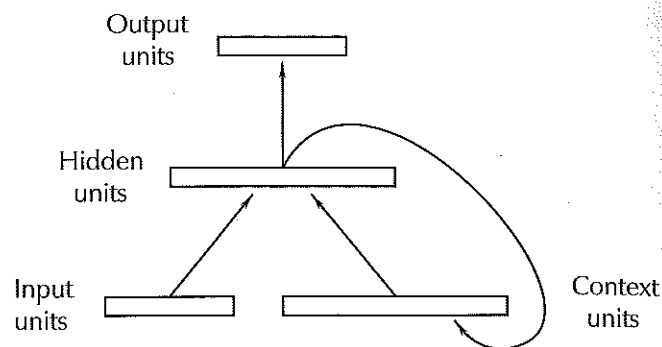


Figure 8.6: A recurrent network of the type employed by Elman (1986) for learning to recognize sequences.

output values. This basic structure, with numerous variations, has been used successfully in producing sequences of phonemes (Jordan, 1986), sequences of movements (Jordan, 1989), sequences of notes in a melody (Todd, 1988), sequences of turns in a simulated ship (Miyata, 1987), and for many other applications. An analogous network for *recognizing* sequences has been used by Elman (1988) for processing sentences one at a time; and another variation has been developed and studied by Mozer (1988). The architecture used by Elman is illustrated in figure 8.6. This network also involves four sets of units: *input units*, in which the sequence to be recognized is presented one element at a time; *context units*, which receive inputs from and send outputs to the hidden units and thus constitute a memory for recent events; *hidden units*, which combine the current input with the memory of past inputs either to name the sequence, to predict the next element of the sequence, or both; and, of course, *output units*.

Another kind of architecture that has received some attention was suggested by Hinton and has been employed by Elman and Zipser (1987), Cottrell, Munro, and Zipser (1987), and many others. It has become part of the standard toolkit of backpropagation. This is the so-called method of autoencoding the pattern set. The basic architecture in this case consists of three layers of units as in the conventional case; however, the input and output layers are identical. The idea is to pass the input through a small number of hidden units and reproduce it over the output units. This requires the hidden units to do a kind of nonlinear, principle-components analysis of the input patterns. In this case, that corresponds to a kind of extraction of critical features. In

many applications, these features turn out to provide a useful compact description of the patterns. Many other architectures are being explored as well. The space of interesting and useful architecture is large and the exploration will continue for many years.

## 2.2 The scaling problem

The scaling problem has received somewhat less attention, although it has clearly emerged as a central problem with backpropagation-like learning procedures. The basic finding has been that difficult problems require many learning trials. For example, it is not unusual to require tens or even hundreds of thousands of pattern presentations to learn moderately difficult problems—that is, problems whose solutions require tens of thousands to a few hundred thousand connections. Large and fast computers are required for such problems, and it is impractical for problems requiring more than a few hundred thousand connections. It is therefore a matter of concern to learn how to speed up the learning so that it can learn more difficult problems in a more reasonable number of exposures. The proposed solutions fall into two basic categories. One line of attack is to improve the learning procedure, either by optimizing the parameters dynamically (that is, change the learning rate systematically during learning), or by using more information in the weight-changing procedure (that is, the so-called second-order backpropagation procedure, in which the second derivatives are also computed). Although some improvements can be attained by these methods, in certain problem domains the basic scaling problem still remains. It seems that the basic problem is that difficult problems require a large number of exemplars, however efficiently each exemplar is used. The other approach grows from viewing *learning* and *evolution* as continuous with one another. On this view, the fact that networks take a long time to learn is to be expected, because we normally compare their behavior to organisms that have long evolutionary histories. Accordingly, the solution is to *start* the systems at places that are as pre-suited as possible for the problem domains to be learned. Shepherd (1989) has argued that such an approach is critical for an appropriate understanding of the phenomena being modeled.

A final approach to the scaling problem is through modularity. Sometimes it is possible to break a problem into smaller subproblems and train subnetworks separately on these. Larger networks can then be assembled from those pretrained modules to solve the original problem. An advantage of the connectionist approach in this regard is

that the preliminary training need only be approximately right. A final round of training can be used after assembly to learn the interfaces among the modules.

### 2.3 The generalization problem

One final aspect of learning that has been looked at is the nature of generalization. It is clear that the most important aspect of networks is not that they learn a set of mappings but that they learn the function implicit in the exemplars under study in such a way that they respond properly to cases not yet observed. Although there are many examples of successful generalization (e.g., the learning of spelling-to-phoneme mappings in Sejnowski and Rosenberg's NETtalk, 1987), there are a number of cases in which the networks do not generalize correctly (see Denker et al. 1987). One simple way to understand this is to note that for most problems there are enough degrees of freedom in the network that there are a large number of genuinely different solutions to the problems—each of which constitutes a different way of generalizing to unseen patterns. Clearly not all of these can be correct.

Weigend and I have proposed an hypothesis that shows some promise in promoting better generalization (Weigend and Rumelhart 1991). The basic idea is this: the problem of generalization is essentially the induction problem. Given a set of observations, what is the appropriate principle that applies to all cases? Note that the network at any point in time can be viewed as a specification of an inductive hypothesis. Our proposal is that we follow a version of Occam's razor and select the *simplest, most robust* network that is consistent with the observations made. The assumption of robustness is simply an embodiment of a kind of continuity assumption that small variations in the input pattern should have little effect on the output or on the performance of the system. The simplicity assumption is simply to choose—of all networks that correctly account for the input data—the net with the fewest hidden units, the fewest connections, the most symmetries among the weights, and so on. We have formalized this procedure and modified the backpropagation learning procedure so that it prefers simple, robust networks, and, all things being equal, will select those networks. In many cases it turns out that these are just the networks that do the best job generalizing.

## Connectionist Modeling: Neural Computation / Mental Connections

9

Paul Smolensky

1989

In the past few years the approach to cognitive science and artificial intelligence known as *connectionist modeling* has dramatically increased its influence. Connectionist systems are large networks of extremely simple computational units, massively interconnected and running in parallel. Each unit or processor has a numerical *activation value* which it communicates to other processors along connections of varying strength; the activation value of each processor constantly changes in response to the activity of the processors to which it is connected. The values of some of the units form the input to the system, and the values of other units form the output; the connections between the units determine how input is transformed to output. In connectionist systems, knowledge is encoded not in symbolic structures but rather in the pattern of numerical strengths of the connections between units.

The goal of connectionist research is to model both lower-level perceptual processes and such higher-level processes as object recognition, problem solving, planning, and language understanding. The rapidly growing collection of connectionist systems includes models of the following cognitive phenomena:

- speech perception,
- visual recognition of figures in the "origami world",
- development of specialized feature detectors,
- amnesia,
- language parsing and generation,
- aphasia,
- discovering binary encodings,
- dynamic programming of massively parallel networks,
- acquisition of English past tense morphophonology from examples,