.op
.po 13
.mt 7
.mb 9

Chapter V

Connectionism as an Hermeneutic Model of Mind

> "In my view, the initial intoxification with cognitive science was based on a shrewd hunch: that human thought would turn out to resemble in significant respects the operation of the computer, and particularly the electronic serial digital computer ...
>
> [However] one of the chief results of the last few decades has been to call into question the extent to which higher human thought processes - those which we might consider most distinctively human - can be adequately approached in terms of this particular computational model.
>
> ... the kind of systematic, logical, rational view of human cognition that pervaded the early literature of cognitive science does not adequately describe much of human thought and behavior. [Gardner 85: 43-44]

Preceding chapters present a case for cognitive oriented anthropology in particular that echoes Gardner's observations on the state of cognitive science in general.

Overall, current cognitive theory (inside and outside anthropology) is experiencing a resurgence of positions that are characterized as interpretivist, hermeneutic, or monistic - partially in reaction to the perceived failure of formalist (computational) theories and models. Increasingly, mind is seen as a kind of socio-cultural construction. [Coulter 83, Winograd 86] In many cases this involves a rejection of the idea of building a model (particularly one that can be realized on a computer) of mind as well. It seems to me that this reaction is too extreme.

Whether they prove to be fundamentally sound or unsound, the computational models of mind that have been developed and explored to date must receive credit for a rapid expansion in our understanding of cognition. Building and testing a model to illustrate a theory quickly exposes the flaws and strengths of that theory. [Ideally] years of finely reasoned debate can be resolved in an instant upon the success or failure of an implemented model. For these, and other, reasons it would seem premature to abandon attempts at model building regardless of the strength of one's orientation towards the hermeneutic.

A number of alternative foundations (or at least metaphors) for constructing such models can be found in the

literature of AI. [See West 88 for a synopsis of some of the leading alternatives.] In this chapter one specific alternative will be presented and extended in preparation for making the case that it offers potential value to the study of cognition in relation to culture.

Connectionism - Neural Networks

There is general agreement that the human brain is integral to the expression of human mind. The architecture of the brain is radically different from that of a digital computer, and it seems capable of supporting exactly those cognitive functions that are difficult if not impossible to realize on a computer. Why not, then, build a computer that functions analogously to the human brain? This approach was in fact undertaken in the early years of computer science [Pitts 47, McCulloch 43, Rosemblatt 62] but was almost totally abandoned for ten years, in large part because of the criticism of Minsky and Papert. [Minsky 69]

Recently, interest in this approach has been revived, the criticisms of Minsky and Papert dismissed as having been directed against a straw-man version of neural networks, and some initial successes have been experienced. [Rumelhart 87] Several labels, each reflecting a different

perspective on the general problem, have been attached to this approach, including: neural computing, parallel distributed processing, and connectionism. Connectionism will be the label employed in the following discussion and most of the subtle differences reflected in the various labels will be ignored.

A computer implementation of a connectionist model consists of hundreds or thousands of individual nodes, each of which is analogous to a neuron in the human brain.[1] Each node is connected to a multiplicity of other nodes just as individual brain neurons are connected to others via axonic and dendritic networks.

The functioning of each node is limited to "firing" (discharging an electrical pulse) or not firing - in essence, acting as a simple switch. Whether or not a neuron fires is a function of its inputs. If the input is a direct stimulus from the outside environment the neuron analog will

1 The fact that computer models contain thousands of nodes while the brain contains millions (and billions of connections) introduces a "scaling" problem. There is no guarantee that insights developed with small scale models will "scale up" and be directly emulated in a full scale environment like the brain. Other aspects of this problem will be noted throughout this chapter.

fire. If the input consists of impulses from other nodes in the network these impulses are summed[2], and if they exceed a threshold value the neuron fires. This process emulates the operation of synapses in the brain. [See Figure 1a]

A complete[3] network consists of multiple nodes arranged in a series of layers, one input layer that receives stimuli from the external world, several "hidden layers," and one output layer in which a pattern of firings represents the network's response to the inputs in a form perceptible to the external world.4 [See Figure 1b]

2 Summation of inputs is not a simple process. Complications derive from signal decay, asynchronous arrival of input signals, back-propogation properties, potential loops, inhibited inputs, lack of simple "layering" in the network architecture, etc. These technial issues will not be directly addressed here except to note that they exist are are areas of ongoing research.

3 A similar caveat must be registered concerning the depiction of the connectionist network in general. What appears here is a major simplification that does not deal with a number of very important, but technical, issues. The thrust of the argument presented in this dissertaion is not dependent on detailed examination of those technical issues and their discussion is deferred to other times and places.

4 Layers are determined by the physical wiring patterns employed to construct the network. Although an ideal situation might have every node
[continued next page]

Exactly which nodes in the output layer are triggered is a function of which input nodes were stimulated, and more importantly by the various firing thresholds of the inter-connections among nodes in the hidden layers.

The threshold values that must be met or exceeded by the summed inputs to a node are variable and may be adjusted so that weak signals can be amplified, strong inputs weakened, or inputs completely inhibited. This variability in threshold values (usually called connection weights) emulates the synapses of the brain whose variable resistance to the signals passing along dendrites determines whether or not a brain neuron fires.

Connection weights are critical for two reasons. First, they collectively determine how the network will eventually respond to a given set of inputs. Second, they provide the mechanism whereby a network can be modified so that consistent inputs will produce consistent outputs. For example, if a set of nodes at the input layer are triggered by a specific stimulus - say a digitized representation of a

[continuation] connected to every other node such a scheme is not physically realizable when the number of nodes gets large. There are technical considerations that determine the physical architecture of neural networks that will not be discussed here. [See Hecht-Nielsen 88 for a description of common types.] For our purposes the simple arrangement is Figure 1b is sufficient.

person's face - and the output layer consists of a lamp labelled with a person's name, it is possible to adjust the connection weights of the network in such a manner that every time Sara's face is input then the lamp next to Sara's name will be lit.

Neural networks can be "trained" to produce specific output patterns in response to a range of specific input patterns by providing feedback[5] which causes alteration in the connection weights which, in turn, causes the output pattern to change. Given the correct algorithm for adjusting connection weights, appropriate feedback, and numerous iterations the network will eventually "learn" the connection weights that produce desired outputs to given stimuli. [See McClelland 87 Vol I, Chapters 7 & 8 for a discussion of algorithms, usually called "learning rules," developed to date.]

A famous example that illustrates how adjustment of connection weights modifies a network's ability to produce consistent outputs is NETtalk, developed by Sejnowski and Rosenberg, which accepts strings of letters as input and

5 Again, the specific and technical details of a feedback mechanism are difficult to conceive and implement. They will depend upon both the architecture on which they will be implemented and the problems to which they will be applied.

produces a series of audible phonemes. Initially the output is random babble but eventually the system begins to correlate the character inputs with appropriate phonemic outputs, and understandable words are formed. The system is able to achieve a 92 percent accuracy level and produce an understandable but somewhat labored verbal rendition (about on par with a child's performance) of a page of printed text. [Sejnowski 86]

Connection weights are not only central to the construction and operation of neural (connectionist) computers but are fundamental to the conception of mind embodied in that type of computer. The best way to illustrate this importance is to consider differences between formalist "representation" and connectionist "distributed representation."

Distributed Representation

As noted in previous chapters a prime tenet of formalist conceptions of mind involves the representation (as symbols or tokens) of the external world in the mind of the perceiver of that world. Cognition is taken to be the manipulation of those symbols according to a set of formal rules. Symbols, in this scheme, are discreet entities.

This conception is consistent with the operation of conventional serial computers which also manipulate discrete symbols according to a set of rules (a program). In a conventional computer specific symbols are represented by digital values stored in specific physical locations comprising the computer's "memory."

In a neural computer, however, any given symbol is represented by the collectivity of connection weights in the computer. Each symbol is "distributed" across the entire system rather than associated with a particular component of that system.

A somewhat basic illustration of how the differences in representation are reflected in behavioral characteristics of the system is a comparison of the consequences that accrue when part of a system "memory' is destroyed. Eliminate a storage location in a conventional computer and you have eliminated any trace of the symbol stored in that location. In a neural network, however, the entire collectivity of connection weights would need to be destroyed before the symbol would be lost. Elimination of any single connection weight can be offset by the network simply by adjusting the remaining weights in compensation.

Research has shown, in fact, that any given set of weights can be recovered if a subset of those weights - less than a third - remain and are "fixed" for the period of recovery. [Hinton 84] This capability of neural networks (in direct contrast to those of conventional computers) is reminiscent of the human brain where it is also possible to destroy portions of the architecture without causing a complete loss of its "contents."

It could be argued that connectionist systems do not contain any representations at all, that what is stored in the system is a "representational complement" of a symbol rather than a symbol as such. Again, a comparison between conventional computers and neural computers will illustrate this difference. In a conventional computer an input pattern is compared with and matched to a stored pattern when that system is said to "recognize" the input pattern. In a neural computer no such match takes place. Instead the signals generated by the input pattern are "channeled" along pathways in the network until they "settle" in those nodes whose discharges constitute the output pattern appropriately responsive to the input pattern.

Technical descriptions of distributed representation, learning functions, and pattern recognition capabilities of neural networks involve various levels of mathematical

description, from simple linear algebra, to probability, to matrix algebra, and into non-linear calculus. [For an overview see Grossberg 88, Carpenter 87, Feldman 82, and McClelland and Rumelhart 86] Because mathematics can obscure as well as reveal, the metaphor of a landscape was introduced to facilitate the understanding of distributed representation and of the operation of neural networks.

> "Consider a countryside laden with with hills and valleys. In the valleys lie lakes. If you pour a bucket of water on a hill, it flows down the hill into one of the lakes. No matter where you pour the water, it will eventually come to a place to rest; the system of mountains, lakes, and flowing water will eventually reach a stable state. And just as there are many mountains and lakes, there are many different stable states the system can go to.
>
> ...neural nets have contours like the hills and valleys in a countryside; they also have stable states. [Hopfield 86: 27-28]

The landscape metaphor is employed in two distinct ways; one "passive" in which the topology of the network is said to represent a symbol; and one "active" in which it describes the operation of the network and is actually representing a procedural complement of a concept. Although this distinction is very real, it is seldom made explicit in connectionist literature. Additional discussion of this point will be presented in the next section.

Whether a given representation in a neural network is direct or complementary it is always holistic - that is, it is realized by the entire collectivity of connection weights rather than an individual weight or node or small set of weights or nodes.

Landscapes, Learning, and Complementary Representation

"Fresh from the factory" a neural network is a tabula rasa representative of nothing and capable of performing no tasks. So too is a conventional computer. The latter must be programmed before it can perform useful functions. The former must be "trained." "Training" can be illustrated with a simple analogy.

Imagine, for example, a simple piece of cloth. As it lies flat upon a table it represents nothing (not even the concept of nothing or "null"). Pick it up and drop it and it will assume some convoluted shape upon the table top; a shape that can be thought of as a representation of the input (the process of picking it up and dropping it).

Analogs in a neural network, 1) the cloth :: the collectivity of connection weights; 2) the pick-up-and-drop process :: stimulation of certain nodes in the input layer of the network, and 3) the at-rest state of the cloth :: the firing of nodes in the output-layer of the network.

If we pick up the cloth and drop it again there is a negligible chance that it will assume the same shape upon landing on the table top. This is true even if we recreate with exactitude the conditions of the original maneuver. Because the connection weights in a network can be adjusted (while the fibers in the cloth cannot) it is possible to "train" the network to recreate with precision the original configuration even though it is highly unlikely for the cloth to re-assume the original shape.

Network "training" might consist in providing feedback to the effect that the second "output" varied from the first configuration and was therefore "wrong." (In some instances the degree of error might also be part of the feedback.) A simple algorithm[6] is applied to each connection weight in the network to adjust that weight up or down and the "training stimulus" is re-applied. This process repeats

6 A number of these algorithms (usually called training rules) have been developed. Hebb [49] proposed an averaging rule that has been modified and generalized as a general delta (GD) rule [Rumelhart 86] and a number of other rules have

been proposed and used as suitable to specific situations. Although some rules offer advantages in certain situations there is reason to believe that the GD rule is sufficient even if not optimally efficient to provide the desired "learning" behavior for a neural network.

until the network reproduces the original configuration. If the weight adjusting algorithms (training rules) were correctly selected, the time required to "learn" (which is dependent upon the overall scale of the network - number of nodes and connections) a given set of behaviors will be significantly reduced.

Network configurations (the collectivity of connection weights) are usually called E-surfaces.4 During "training" the E-surface of a network is equated with a representation of the input stimulus. It is possible, however, to train a given network to "recognize" multiple inputs. The result is an E-surface that increases in complexity - its convolutions simultaneously "representing" multiple inputs.

4 "E" for energy, energy being the analog of altitude in the landscape metaphor. An E-surface is characterized by high and low energy locations (mountains and valleys). When a network is stimulated with an input the E-surface is disturbed, but the operation of connection weights ensures that it will return to a point of stability with the energy introduced by the input having been channeled to one of the low-energy locales connected to appropriate output-layer nodes. [See McClelland and Rumelhart 86 for the

technical and mathematical description of this process.]

At the point when a single E-surface begins to simultaneously represent multiple (and potentially contradictory, like "yes" and "no") inputs, severe damage is being done to our normal conception of representation; e.g., that "X" stands for "Y." As noted earlier, however, there is a second sense in which "representation" is used in connectionism, a sense where the term more properly refers to a procedural complement of common-sense representation.

Specifically, the E-surface of the network represents not the input stimulus directly but the "channeling process" which assures that presentation of that stimulus to the network will result in the generation by the network of an output deemed appropriate for the stimulus, e.g., producing the correct name when presented with a photograph of a face. It is this second sense (or perhaps contra-sense) of representation that gives rise to the landscape metaphor and the metaphorical flow of water over and down that landscape until it reaches the lowest possible point, an equally metaphorical lake.

The set of inputs that a network has "learned to recognize" and the repetoire of "meaningful" outputs that the network produces can be considered as the external environment of that network. Internally, the topology (configuration of connection weights, landscape) of the network is most properly considered as the complementary representation of the network's environment as-a-whole rather than the simultaneous "X stands for Y" type of representation of a set of stimuli and a repertoire of "IF X THEN Y" response patterns.

Network topology is a complementary representation because it "stands for" the process whereby input X results in output Y rather than "standing for" X and Y directly. In terms of the metaphor, the landscape represents the channels whereby water is appropriately transported from an arbitrary point on the surface into one of potentially many lakes rather than the water drop and the lake. Network topology is a representation because it does "stand for" the external environment of the network.

Additional Aspects of Neural Networks

There are six additional aspects of neural networks that are important to future discussions and which need to be briefly introduced.

1) Although neural networks are realizable in terms of hardware the topology of the network is a virtual, not a physical, entity. Changes in the environment are accommodated by the network with changes in its virtual

topology, not with a physical restructuring of its hardware. This feature makes a neural network extremely adaptable. In contrast to an organism which can be seen as a special purpose adaptation to its environment, a neural network is a general purpose, dynamic, adaptation device.

2) Neural networks are intrinsically wholistic. As noted previously, the network topology is complementary to the external environment of the network in the same way that one piece of a puzzle is complementary to another. To illustrate this contention it is useful to look at another biological model, protein enzymes.

> "The key computing attribute of protein enzymes is their folded shape. Recall that recognizing patterns and objects is a difficult problem for computers. Pattern recognition, however, is the main activity of protein enzymes. Their folded shape allows these enzymes to recognize molecular objects on the basis of tactile (touching) interactions reminiscent of the way a key fits into a lock. The switching action of the enzyme - making or breaking a covalent bond - is secondary to the recognition process. In effect, the enzyme is both an intelligent and an evolvable switch.
>
> Any computational function that can be implemented using conventional switching elements (such as the McCulloch-Pitts formal neurons) can be implemented using tactilizing processors, and, in general, much more efficiently. All conventional switches do is recognize simple patterns (such as 11 or 10). Recognizing complex patterns requires networks of many simple switches, whereas tactilizing

> processors are capable of recognizing complex spatio-temporal patterns by themselves. [Conrad 87: 13-14]

Although, strictly speaking, the topology of a neural network is not the "recognition" of an external pattern it is congruent with the external environment. And it is congruent on a whole-to-a-whole basis. This distinction is important because it differentiates a neural network topology from a straightforward simulacrum. [This biological example of wholism will be referred to again later in this chapter.]

3) A neural network can function only in response to an environment that is regular and recurrent. As noted earlier the topology of a network is not "pre-programmed" (nor is there any known way to perform such programming), it "learns" its configuration through multiple iterations of the adjustment process. It is easy to see that, if the inputs do not recur, iteration is impossible, learning cannot occur, and no topology is configured.

Regularity is a need that is not quite as apparent a need as recurrence, but it is closely related. It would be unreasonable to expect that a given stimulus will be presented for each iteration of the learning process in exactly the same way. If the form taken by the input is allowed too much variation, however, the network is

essentially confronted with a non-recurrent environment. It is difficult to precisely define the limits of allowable variation but Figure Two illustrates why some constraints must be placed on variation. In each of the three examples it is reasonable to expect the network to "recognize" as the same object the versions in examples (a) and (b) but not the examples marked (c).

Figure Two: Input Variability

4) Because a network topology is "learned" and because that "learning" is effected in part by feedback (either internally generated or externally supplied),it is necessary

for that feedback to be consistent. This requirement is seemingly too obvious to require mention; it is immediately apparent that inconsistent feedback would prevent the network from ever settling on an appropriate configuration. However, the point will become important in later discussions when the performance expected of networks is transported to the natural realm.

If the human mind, for example, is realized by a network of the type we are discussing, we will need to be able to account for the fact that it manages despite receiving inconsistent feedback. (Consider the inconsistent and often contradictory injunctions leveled on children, for example.)

If a connectionist model is to be robust, the need for consistent feedback must be modified to allow for some variation or for some sort of hierarchical meta-consistency. (An example would be the consistency provided by a moral code that superceded that provided by an injunction to obey laws involved in a logical conflict.)

5) As presented so far, there is no way that a neural network can distinguish between an object and the ground upon which that object exists.

There are no intrinsic differences between one node in a network and any other nor between one connection and any

other. There is no intrinsic bias or filter built into the network that cause it to "pay attention" to a subset of input nodes or to a subset of processing nodes (or connections).[7] Therefore, both ground and object are equal sources of information to the network or, more accurately, ground and object are indistinguishable parts of the same, whole, input.

Differentiation between ground and object is a necessary precondition to explaining observed characteristics of the human mind like "concentration" and "attention." To provide an adequate model of mind (at least human mind), the ability to make ground-object distinctions will need to be developed, ideally by extension of an existing characteristic of neural nets.

6) The whole of a neural network topology can be reconstructed from a subset of the participating connections. Specifically, if a subset of the connections

7 In the case of vision experiments this is not necessarily true. A network may be explicitly trained to recognize features (edges, surfaces, etc.) rather than the whole input pattern. This does not reflect a limitation of the network. The situation arises because researchers are using networks to emulate processes (and consequently the limitations of those processes) as developed for serial computers. [See Marr 82]

in a network are fixed and the remainder of the connections are "scrambled" the network can (and will) re-create the original configuration. Hinton uses the illustration of a room full of students flipping a switch on or off in response to a light on the switch panel. If a certain number of the switches are permanently fixed in the on position and the remainder of the network is randomized then the subsequent actions of the students in response to the lights will result in the restoration of the original configuration.

At first this characteristic of a network seems to contradict, at least in part, the wholism discussed previously. It seems as if a portion of a configuration can "stand for" (represent?) the whole. It must be remembered, however, that the configuration is invoked by a percentage of nodes N1 through Nn, not a subset of nodes (N1, N5, N12 and N103). What is observed is a consequence of distributed representation,[8] not a manifestation of particularism in an otherwise wholistic network.

8 One that is reminiscent of a property of holograms (another example of distributed representation). A hologram in the form of a piece of photographic film can be cut into parts and each part will retain the ability to regenerate the whole image - with a loss of resolution inversely proportional to the area of the fragment.

Neural or connectionist networks as discussed so far provide the foundation for constructing an alternative model of mind but are in and of themselves insufficient. Neural networks provide a model of how cognition (thinking) might occur - it focuses on a process. A model of mind, however, must also include linkages between that process (the neural network) and the context within which the process operates. To accomplish this task an additional element will be proposed for the neural architecture (following section) and the landscape metaphor will be extended to illustrate the functionality of the new element (subsequent section).

Constraint Windows

Current implementations of neural networks provide for an initial assignment of connection weights and a learning algorithm that determines adjustments to those weights as a function of feedback. Initial values may be assigned to the connection weights, no assignments may be made, or completely random assignments may be made.

Assignment of initial weights does not contribute to the capability of the network. In some instances a performance

benefit is derived; a judicious initial assignment can significantly shorten the time (number of iterations) required for the network to "learn" a given task. The fact that a quantitative performance benefit can be derived from an action (initial weight assignments) without making a qualitative difference in the network's capabilities or characteristics will be of potential use and will be discussed later in this section.

The standard architecture of a neural network (nodes, layers, connections, connection weights, feedback, and a learning algorithm) reflects a focus by connectionists on the electrical aspects of the brain to the relative exclusion of all other aspects. This focus has been criticized as reductionistic; it is certainly incomplete.

Neural (connectionist) approaches are modeled after the operation of the human brain (our existence proof of mind).

A quick examination of the brain, however, reveals the fact that neurons and neural connections are affected in subtle (sometimes not so subtle) and complex ways by brain chemistry. The operation of the neural network (and hence its internal configuration, E-surface, or topology) can be enhanced, inhibited, and totally disrupted by alteration of the chemical environment of that network. A more complete model of a neural network (assuming the brain to be such a

network) would therefore allow for the influence of what Bergland [85] calls an "interwoven fabric of chemical threads," in a brain that is a "gland."

Bergland reviews the development of ideas about the brain, both its structure and its place as the seat of the mind, from Plato to the present. He notes the development of the "electrical" paradigm - the notion that the brain is a complex electrical circuit the configuration of which and the passage of electronic signals through which constitutes "thought" - and challenges this paradigm with recent findings in neuro-physiology and neuro-endocrinology.

A key observation for Bergland is the architecture of a synapse, "the grasping claw," that is the central component in the electrical circuitry of the body. Synapses are "circuit breakers" in that there is a discontinuity in every "brain circuit" at the point of a synapse, a discontinuity that must be bridged the way a spark jumps across space in many familiar circumstances. Synaptic closure - ability to complete the circuit by sparking - is a function of brain chemistry. Different chemicals (actually hormonal molecules) inhibit or enhance the ability of the synapse to close.

Although Bergland does not directly refer to connectionism or neural networks it seems clear from his

general arguments that he would find such approaches to be as reductionistic as other computational models of mind. Most likely he would object to the idea that the configuration of a network arises from nothing more than a function of the weighted firings of nodes. He probably would insist that (to be complete) the model would have to include an element that simulated brain hormones and their ability to act as "gatekeepers" determining which neurons will fire and in which circumstances.

In one sense Bergland is proposing nothing more than a neural network the topology of which is determined by chemistry (hormones) in addition to electricity. However, including a mechanism in a neural network that emulates or accounts for brain chemistry (hormones) will have significant consequences for model building. First, it emphasizes the role of pattern recognition (the wholistic pattern recognition of enzymes noted previously):

> "As scientists accept this new paradigm, the primary mechanisms of intelligent thought must be viewed differently. The mind is made pattern dependent and comes to share in the ubiquitous secret of

> evolutionary survival: pattern recognition. [Bergland 85: 108-109]

More radical is a redefinition of the locus of cognition:

> "The mechanisms of the mind are thus released from the conceptual confines of the reductionistic left brain. The

> mechanisms that drive thought are found all over the body and, wherever they live, they function at their highest level by recognizing the molecular patterns of the combination of hormones that modulate thought." [Bergland 85: 109]

As radical as Bergland's ideas may seem, they can be accommodated with a relatively minor change in network architecture - replacement of "thresholds" with a mechanism we will label a "constraint window."

Whether or not a node in a standard neural network will fire is determined by summing the inputs (multiplied with the connection weight assigned to that input connection) to that node. If the sum exceeds a certain threshold[9] then the node will fire. Where a threshold provides only a floor value, a constraint window will provide both floor and ceiling values. A given node will fire only if the summed inputs exceed the floor value but do not exceed the ceiling value. [Figure Three provides a graphical illustration of a threshold (a), several examples of constraint windows (b) (c) (d) and (e), and a revised diagram of the components in a neural network (f).]

9 Thresholds are based on an arbitrary selected range values, usually from 0 to 1. A

threshold is exceeded if the summed input currents exceed the value assigned to the threshold. The range of possible input measures will be referred to as the input value range or value scale.

A constraint window with a relatively broad range that is congruent with the upper end of the value scale [See fn. 9 on preceeding page.] functions identically to a threshold. [Figure 3(b)] How windows defined with other parameters affect the operation of a network can best be understood in terms of the landscape metaphor.

A window with a narrow range centered near the top of the scale has the effect of creating a semi-permanent (until the parameters of the window are changed) "peak" in the network's topology. [Figure 3(c)] One with a narrow range centered near the bottom of the scale creates a "valley." [Figure 3(d)] And, one with a moderate range centered at the mid-point of the scale creates a node that will always (very high probability) fire, one that is locked in the "on" position. Closing the window locks a node into the "off" position.

It is important to remember that although the specific effect of a constraint window is on the performance of an individual node, the general effect is to "pre-dispose" the virtual topology (configuration) of the network as a whole. "Pre-disposition" in the sense used here is identical to the pre-setting of connection weights or the "freezing" of a subset of nodes in order to "recover" a configuration. [See preceding discussion.] No material change is made in the operational characteristics of the network by adding a constraint window, but, as we will see later, it has important effects on the observed behavior of a network.

[This discussion of constraint windows omits numerous important details concerning their implementation and actual operation. Investigation of these details remains to be undertaken and is noted as an area for further research in Chapter Seven.]

Extending the Landscape Metaphor

When constraint windows are added to the architecture a mechanism is provided whereby network topology is influenced by "secondary factors." ("Secondary" does not imply inferior or less significant - it implies "additional.") As discussed in Chapter IV, the primary aspect of the hermeneutic conception of mind requires that cognition be a function of the context in which it takes place as well as of the specifics of an individual cognitive act. Constraint windows provide the mechanism (alluded to by Geertz [Geertz 73:82]) whereby context can be accommodated.

Given the mechanism, how then to describe the context? What are the "secondary factors" that need to be recognized? Bergland's glandular chemistry is one approach to defining a

set of secondary factors. Chemistry, however, is too close to being a mechanism itself to properly be considered as the kind of secondary factor implied (required) by the hermeneutic concept of mind.

At least six broad notions of a secondary factor can be distilled from the hermeneutic concept of mind as discussed in previous chapters. Identifying those notions and relating them to the neural network model as developed so far can best be accomplished by employing the landscape metaphor already introduced.

A natural topology, the surface of a planet for example, does not exist in isolation. It is determined in part by the underlying geology of the planet and by the environment that works upon that surface. It would seem reasonable, therefore, to extend the topology metaphor of neural networks by analogy to a natural topology.

Beginning with the most basic, the six secondary factors are:

> Cellular - [the "planetary core"] Bergland [85] and Conrad [87] discuss in detail the "information processing" capabilities of enzymes and cellular (or sub-cellular) entities. Many of the constraints imposed on thought originate and are influenced by activity at this level. Examples range from basic

> sensitivity to stimulii (a fundamental constraint) to the effects on cognitive activity of an "adrenaline rush."

Organismic - ["tectonic plates"] Maturana and Varela [87] present a carefully developed argument that roots cognition and cognition based behavior in the characteristics of cells and the organisms that are structured collections of those cells. The major portions of their argument are focused on organism level phenemona. Winograd and Flores [86] extend the arguments of Maturana and Varela specifically to the modeling of cognition in AI.

In addition to demonstrating how organismic organization affects cognition and behavior, Maturana and Varela show the arbitrary nature of our classification of behavior into "intelligent, cognitive, and aware" on one hand and "instinctive, stimulus-response, and non-intelligent" on the other.

Although specifics of their arguments are open to challenge, their general argument - that so-called higher functions like cognition are constrained by the structure of the organism (which is reflective of both its ontogeny and phylogeny) in which they occur - is readily accepted.

Sensual - ["planetary crust"] This is the level of the five known senses and the perception data that they feed into the mind. It is the acknowledged starting point for most models of cognition. It is also the focus of some of the major philosophical debates on the nature of mind and its relation to the external environment.

For our purposes two aspects of this level are of special relevance; its apparent complexity and the large percentage of it which escapes our "attention." One major focus of AI research can be characterized as the emulation of the sensory capabilities of human beings; vision research being the most prominent example.

Experience in vision research confirms the complexity involved in simple sensory recognition of external stimuli. This complexity is evident even when the problem is restricted to that subset of external stimuli paid conscious attention. For example, the complexity involved in the problem of sensing the keyboard and monitor of the computer that is used to write this paragraph pales in comparison to the vast sum of sensation that is being "ignored" - voices out the window, my son practicing piano downstairs, the

feel of the chair against my thighs, and the objects intruding via peripheral vision.

The fact that I can enumerate (partly) the senory information received outside of my attention sphere indicates that it is present whether or not I pay attention to them. It is reasonable to assume that "background sensations" continue to play a role in the shaping of cognition even though they lack "attention." This point will be discussed further in the next section.

Cultural - ["geography"] This is the level of central importance to this thesis. The heart of the hermeneutic argument involves the social construction of meaning, the cultural parameters of cognition, and the behavior-symbol-cognition associations that are empirically evident. Details of this level have been and will be presented throughout the thesis.

Habitual - ["landscape"] Individual performance and individual thought exhibit an individual level of consistency - habits. Like the cultural level, the habitual level is implicit in all of the discussions herein. Habits are important not only for their proactive role in shaping cognition but also as the primary source of individual variation. Although

variation is introduced at every level (no two people have identical DNA, organisms share phylogeny but not ontogeny, etc.) the most obvious variability occurs at this level.

Analytic - ["architecture"] Just as the man-made environment exists as a thin veneer on the natural landscape, the analytic layer represents that thin veneer of cognitive activities that are popularly known as "thinking." This is the realm of "book learning"

> and of deliberative thought. It is also that realm that has been the focus of early AI efforts (Simon, Newell, Minsky, et. al.), mainstream linguistics, and cognitive anthropology. From the hermeneutic perspective the chief contribution of those efforts has been in demonstrating how thin indeed is the veneer of analytic thought.

Two general points must be raised regarding the extended metaphor and its proposed levels. First, the relationship among levels is not deterministic nor is it hierarchical. All levels operate simultaneously. The point of relating secondary factors in terms of the landscape metaphor employed in explaining neural networks was to show the essential operational unity of each factor in shaping

network performance. However convenient it may be to speak of each influence as if it were isolated from the others, such a perspective is illusory.

Second, the use of metaphorical levels to illustrate multiple axes of influence on cognition should not be confused with "stratigraphic" models as defined and criticized by Geertz. [Geertz 73: 37-45] Specific differences include the existence of individual variation at every level of the metaphor, the lack of any hierarchical relationship of levels, and the fact that no level is "explainable" in terms of "lower" levels. Instead the metaphor is intended to illustrate the manner in which a variety of essentially independent influences can operate simultaneously to shape the empirical features of a specific phenomenon.

The essential components have now been introduced and the remaining task is to assemble them into a sensible model.

Model of Mind

Presentation of the model will consist in part of a summary of preceding sections, in part of a description of what might be called architectural attributes, and in part of a description of operational characteristics.

The physical architecture underlying the model is a densely connected neural network consisting of layered nodes, each node capable of receiving signals from and sending signals to an arbitrary number of other nodes. Whether the input signals are sufficient to cause a node to emit a signal is a function of a thresholded sum and the operation of a constraint window. Except for the constraint window this architecture is a standard neural network.

More important than the details of the physical architecture is the capability of that architecture to perform as a generator of a virtual topology which is a wholistic and simultaneous representation of an input pattern and a representational complement of that pattern.

A representational complement is the set of potentials that assure a given input pattern will yield an appropriate output pattern.

In metaphorical terms the mechanism can be viewed as a generator of a topological surface that constantly varies as a function of received inputs - much the same way that a flag constantly assumes a different topology under the influence of a breeze. Network response need not be passive and random however. Given regularity of input and the presence of topological constraints the virtual surface can

exhibit predictable behavior just as a sail will form an airfoil given a steady breeze and appropriate line tension.

Addition of constraint windows to a typical neural architecture was prompted by the desire to account for additional influences (e.g., chemical-hormonal) which would allow retention of the significance of the brain as the "seat of cognition" without the need to impose a reductionist separation of the brain from its immediate milieu or the environment as a whole.

A major side effect of the constraint window architecture is the inability to differentiate, in principle, between electro-chemical activity in the brain, in the central nervous system, or in the sensory organs. Further, it could be argued that the pattern recognition capabilities of the architecture differ from those of an enzyme (or enzymatic tactile processor [Conrad 87]) only by virtue of being general rather than special purpose. In essence this results in a mind whose dimensions are conterminous with those of the body with which it is associated. This attribute of the expanded neural architecture is consistent with the concept of mind presented by Maturana and Varela as well as Bergland. [See above.]

There are five operational characteristics of the model that need to be introduced and that will complete the description of the model.

1) Existing neural computers are operated in fixed intervals and address specific problems. While this is satisfactory for experimentation it should be made clear that a necessary operational characteristic of the proposed model is its ability to function in a continuous mode. Just as the human mind is always "on" so too must be the neural network. A corollary to this characteristic is the need to accommodate to a continually changing (though regular) pattern of inputs. (Neither of these characteristics are unrealizable, but both face some definite implementational difficulties.)

2) As depicted the model will be continuously responding to a wide range of inputs (effectively the entire sensory and physiological environment). An absolute

requirement, if the network is to function at all, is that the environment exhibit regularity. It is likely, however, that there will be differential levels of regularity. In human terms: inputs from autonomous nervous system functions like breathing will be far more regular than inputs from the optic nerve.

A consequence of differential regularity in the environment is (apparently) a similar differential in the performance of the network. In architectural terms, there are no attributes that would allow this differentiation since all nodes, connections, connection weights, and constraint windows are alike and operate on identical principles. The virtual topology, however, can exhibit this characteristic in the sense that some "features" of that topology (hills and valleys) will be more persistent than others. This aspect was illustrated by the extended metaphor introduced earlier.

When operating a neural network is constantly responding to stimuli. In descending order from the analytic each level provides input that is increasingly regular (consistent). Inputs from the analytic level are highly variable (change with great frequency) while those at the cellular level occur with evolutionary slowness. When this differential in variability is realized in terms of the extended network architecture, it has the effect of providing a high-regularity "ground" against which low-regularity objects can be distinguished.

Differential regularity has two important side effects. First, high-regularity inputs - once established - have potential equivalence to "pre-setting" connection weights in

standard neural networks. High-regularity inputs could (need not, but could) enhance the networks performance when "learning" low-regularity input patterns. Second, differential regularity offers a partial explanation of the "attention" phenomenon - the tendency to focus conscious awareness on the least regular of inputs.

3) Any given configuration of the network can be invoked (or maintained) by "locking" a sufficiently large subset of the nodes participating in that configuration to their appropriate states. [See above.] If a given subset of inputs is of sufficiently high-regularity that they have the effect of "locking" the corresponding set of network nodes then that subset of inputs can be considered as a functional replacement of the greater pattern. This could be true for a specific case or in general. Subsets of this type will be important in the next chapter and will be labelled "kernels."

4) The model provides for dimensional parallelism in the sense that all six of the metaphorical sources of inputs operate on the network simultaneously. Existing models of neural networks essentially simulate the operation of the

focused consciousness of mind - i.e., they focus on a specific set of inputs for a particular problem. A network of the type proposed would place conscious processing inputs

in a context of non-conscious inputs, providing a more accurate simulation of human cognitive processing.

5) By virtue of being "housed in" (or conterminous with) a physical body, mind - as defined here - encompasses a mechanism for interacting with the external environment. It is possible to modify that environment and in so doing modify the inputs from that environment to either increase or decrease their regularity. This operational characteristic will also receive significant attention in the next chapter.

The preceding model of mind is derived from existing neural network models but makes significantly greater claims for their capabilities - based mostly on the proposed addition of constraint window mechanisms to those models - and is therefore vulnerable to criticism until such time as appropriate networks can be built and demonstrated.

"Proving" a model of the type proposed was not the intent. The point has been to present a model that is, in principle, conceptually feasible even if not yet realizable. Further it is argued that such a model is compatible with the demands of the hermeneutic conception of mind and that it could be used to support hermeneutic arguments in the same manner as conventional computer models have been used to support the formalist position vis-a-vis mind.

This chapter presented the foundation and the technical components of the argument. The next will address the utility of the model in explaining or understanding anthropological issues.