Reflections in Silicon: Artificial and Natural Neural Networks

by

Rick Willard Tanney, B.A., B.S., M.Div., MSCIS

Report

Presented to the Faculty of the Graduate School

of The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Master of Arts

The University of Texas at Austin

May 1997

Reflections in Silicon: Artificial and Natural Neural Networks

APPROVED BY

SUPERVISING COMMITTEE

Dedicated to the Old English scribe, Gerefa, who recorded this witness speaking in exasperation at a trial:

Ic gecende be ðam ðe ic cuðe;
se ðe bette cunne gecyðe his mare.[1]

1 Old English : "I have set forth [the subject] as best I could; let him who knows [it] better make more of it known." *Gerefa*. Baugh, Albert C. *A Literary History of England*. New York: Appleton-Century-Crofts, 1967, 36.

Preface: A Look Before the Leap

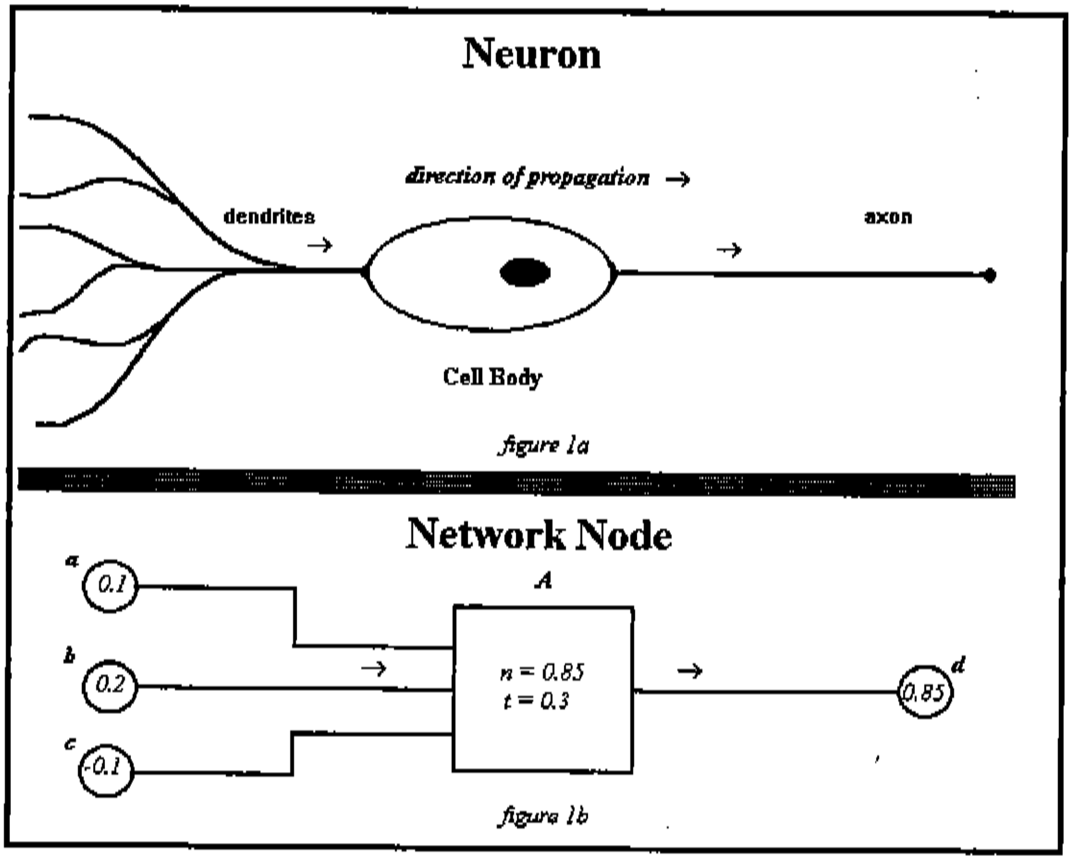

Figure 1. Comparing Neurons and Network Nodes

The basic function of the neuron is this (figure 1a. Neuron)[2]:

1. Input signals arrive from the axons of other neurons through the dendrites

    a. excitatory signals increase the likelihood of activation

    b. inhibitory signals decrease the likelihood of activation

---

2. Rahmann, H., and M. Rahmann. *The Neurobiological Basis of Memory and Behavior*. New York: Springer-Verlag. 1992, 3.

2. Signals are summed and compared to the activation threshold

3. If the activation threshold is reached

then an output signal is sent through the axon to the dendrites of other neurons

else the activation level of the neuron decreases

The basic computer neural network neuron simulation is this (figure 1b Network Node):

1. Input signals arrive from other network nodes through the program algorithm

   a. excitatory signals increase the likelihood of activation

   b. inhibitory signals decrease the likelihood of activation

2. Signals are summed and compared to the activation threshold

3. If the activation threshold is reached

then an output signal is sent through the program algorithm

else the activation level of the node decreases

Many factors affect the flow of activation through the nervous system:

1. Placement in the structure being activated

2. Chemical environment

Factors affecting the flow of activation through a computer neural network are less complex, and include only the algorithm used for propagation, and the data set fed to the system. The basic operation of the individual nodes or neurons, however, seems to be very simple: sum inputs, if the sum is greater than some threshold value, propagate a signal.

This basic operation goes on continuously in the brain. Significant variations occur in the structures active at any given moment. Much text has been dedicated to "brain states," but this term is applicable in a very limited sense. Physically, the only time a brain can be said to be in an actual state is when it is dead. Otherwise, there is a continuous flow. The brain does not maintain a static state while it waits for more input. Nothing is done in the brain by achieving some static arrangement, but only by moving through structures which are themselves constantly subject to modification. Heraclitus' dictum *vis à vis* the river of time applies equally to the brain: you can't progress through the same brain twice. Many general structures remain relatively the same over time, but all are constantly subject to minor modifications. The rate of attrition from disuse of language demonstrates that, without rehearsal, knowledge fades.[3] [4] [5] The brain does not permanently encode its contents, but relies on occasional refreshment or reinforcement. Long-term memory implies long-term reinforcement. The reinforcement need not be from external sources. The unconscious seems to keep busy reinforcing structures by training the brain's neural network through dreams.[6]

3. Brown, Douglas H. *Principles of Language Learning and Teaching, 2d ed.* Englewood Cliffs, NJ: Prentice Hall Regents, 1987, 68-70.
4. Rahmann. *Ibid.* 220.
5. Klimesch, Wolfgang. *The Structure of Long-Term Memory.* Hillsdale: Lawrence Erlbaum Associates, 1994, 12.
6. Fuster, Joaquin M. *Memory in the Cerebral Cortex*: An Empirical Approach to Neural Networks in the

In a computer neural network, the nodes are trained on a given set of data, then frozen while the system is in use doing what it is intended to do. The dynamism which is constant in the brain is only present in the initial training phase of a neural network. Further training may be done, but for the most part, the net is frozen for testing. This paper will focus on the philosophical questions involved in comparing brain function with neural networks, but first I must introduce the basic structures involved in both brain activity and the related neural network paradigms. Both areas are currently in a dynamic state of flux (to use "state" loosely). New brain studies are revealing faults in recent theories of brain function as well as more detailed data regarding how the structures in the brain interact. Computer neural networks are constantly undergoing refinement and facing new challenges. Does knowledge in one of these systems imply similar knowledge in the other? We will first look at the current view of brain function, then the current state of neural networks, before plunging into the details of philosophical arguments about connectionist versus symbolic paradigms, dynamism versus states, single unit versus functional structure, imposed structure versus self-organization, local versus global representation, etc.

Cognitive science is truly a multi-disciplinary field: the philosophical problems are not resolvable by reason alone: reference must be made to the physical truths made manifest by on-going research. The computerized approach will be somewhat isolated

---

Human and Nonhuman Primate. Cambridge: The MIT Press, 1995, 284--288.

in its uses and implications for humanity without some correlation to the philosophical questions relating to *What is man? What is thought? What is consciousness?*. Psychology and linguistics are also deeply engaged in this exploration of *what is?*, *what happens?*, and *how do we value it?* Especially psychology has a great stake in the outcomes of research into the workings of the brain. Who is closer to the truth, Pavlov or Freud, Freud or Jung, Jung or Piaget? The investigation of neural networks is relevant to all these areas. If computer neural networks can be shown to be in a meaningful way equivalent to the neural network of the brain, these investigations could lead to dramatic progress in many significant areas of human daily life, especially in how we think about *how we think*.

The word "model" has a great many uses. In logic it is: "An interpretation I is called a *model* of Γ iff it satisfies Γ. Γ is *satisfiable* iff it has a model."[7] We use it in everyday language to mean a paradigm, as in "a model citizen." It refers to small planning versions of a building. "Model" has many common uses. Consider:

> *I met a girl at a party. She told me she was a model.*
> *I asked, "What scale?"*
> *"Union," she sighed.*[8]

There are scientific and literary models, economic and artist's models, models of trains, planes and automobiles. I will use the word "model" many times in this paper. I mean a particular sense when I use it. This definition is no substitute for the

7. Causey, Robert L. *Logic, Sets, and Recursion*. Boston. Jones and Bartlett Publishers, 1994, 298. "Γ" is a set of Predicate Calculus *(PC)* sentences, and I is a *(PC)* interpretation.
8. Tanney, Rick W. "Toward a Connectionist Lexicon and the Possibility of Disambiguation." Unpublished.

function of either an artificial neural network or of a natural neural network (the nervous system), but I think it may help to link those concepts in both types of network that we seek to elucidate and understand. When I use the word "model", I mean a structure which includes three associated parts: 1) *what is*, 2) *how it works*, and 3) *what value it has*. When I speak of *what is*, I include what physically or psychically exists, to include both existential properties and actions associated with the object or concept. This means our evaluation (conscious or not) of whether the thing in question has physical existence and what its physical properties are, or whether it has a conceptual existence. "Cat" has a physical existence; "addition" exists only as a concept. *What is*, for "cat", will include "a physical presence which is a furry, quick, short mammal with (usually) a long tail and whiskers"; for "addition", it will include "a process by which things are serially joined together." For "cat," this includes such things as jumping, purring, and eating birds; for "addition," it includes combining two elements to form a sum. When I speak of *how it works*, I include what causal relations apply to *what is*. When the cat is stroked, it purrs, that is, *stroking* causes *purring*; when it wants to catch a bird to eat, it jumps -- this exemplifies two causal relations -- *wanting* causes *jumping* which results in *catching*. *What value it has* includes what our interest in it is, its survival value, pleasure value, its relative desirability or acceptability. For "cat", this includes the pleasure I feel when I stroke it and it purrs. For "addition", it includes the practical value in being able to join many things together to be treated as a group rather than going through the tedium of treating each element individually. These three elements combined comprise all relevant factors I can think of that are needed to form a working black box which can process incoming data relating to the thing we have in mind and output a meaningful result. I call this the "Model model."

Science is developing a model of brain function which has implications for our philosophical models of thought, perception and knowledge. Neural network models can be tools for manipulating and evaluating these models. If we can develop valid working models which are essentially the same in both areas, we may be able to use neural networks for many practical as well as philosophical investigations. Natural neural networks are the basis not only of thought, but of memory -- short, medium, and long-term -- which is necessary for our sense of identity. Through artificial neural network simulation, we may discover more about who we are, and how we are aware of who we are.

Reflections in Silicon: Artificial and Natural Neural Networks

by

Rick Willard Tanney, M.A.
The University of Texas at Austin, 1996
SUPERVISOR: Robert L. Causey

Abstract

The concept of neural networks as serious analogs for human brain function is very appealing, but controversy continues over how significant they are, and in what contexts they may be applicable. They are of serious interest for philosophers for their implications for theories of mind, for issues involving thought, perception and knowledge. There is serious doubt about whether an electronic mechanism, no matter how elaborate, can meaningfully mimic the functions of a complex electro-biochemically active organ. Proponents believe that neural networks can perform cognitive tasks at a functional level, and that we may be able to learn much about our own cerebral activities from their study. I conclude that neural networks are an important and essential model for mental function, but that the neuronal circuitry alone is not sufficient to explain motivation or overall system guidance.

Table of Contents

List of Tables Page

List of Figures Page

Introduction

The brain is amazingly complex, yet scientists have been making great strides in discovering its workings.[1] To a great extent, brain functions seem to be based on a few simple principles and procedures involving neurons which are made powerful through the structure of their synaptic connections. New techniques are constantly being developed to allow an ever finer-grained examination of the brain at work. The functions of smaller and smaller portions of the brain's rich structure are being defined with increasing clarity. The black boxes are getting ever smaller. Effective models are being developed at many functional levels, and the contribution of each level to overall function is being integrated into an overall picture. The picture is far from complete, and we will have scant space in this paper for a detailed view of even what is currently known. I will try to paint a broad enough picture not to get bogged down in unnecessary detail, but clear enough to indicate the significance of each level and how it applies to questions of neural network simulation and philosophical implications for theories of mind.

Neural networks as brain simulators are of relatively recent origin, having begun with the pioneering work of McCulloch and Pitts in 1943.[2] There are many forms of neural network, the primary ones being back-propagating networks and self-organizing maps. I will try to paint a broad picture, and focus on the pertinent

1. The 1990's have been designated the "Decade of the Brain."
2. Haykin, Simon. Neural Networks. New York. Macmillan College Publishing Company, 1994, 36.

implications. As computers become faster and increase in capacity, simulations will become more and more effective, and come closer to modeling the real neural networks of the brain. Their range is currently very limited compared to the vast capacity of the brain, megabytes compared with multi-gigabytes. But it is the emulation of structure which is most important, and which is most controversial. Can a mere electronic device act like a very complex electro-biochemical organ? I hope to lay before you a lucid discussion of the issues as they currently stand.

What can we hope to explain using neural networks? Learning and memory -- how do we learn? How do we forget what we've learned? Knowledge representation -- knowledge is not stored within the neurons, but in the structure which connects and activates them. Can we use this structural approach to good advantage elsewhere? What is the relation of the real world to our perception of it? Will our concept of the real world change once we discover *how* we understand it internally? Humor. Puns. Irony and ambiguity. I believe that humor is the result of the overlap of models stored in the neural network. Volition, motivation. What are the mental mechanisms which drive us to do what we do? Spoonerisms. Language -- it's not straightforward symbol processing, so how do we coordinate printed words, their sounds, and our internal representation of what they mean in the real world? *Déjà vu, Presque vu, Jamais vu* -- I've done this before; I almost saw something; I feel like I've never seen this before (although

I see it every day). Intuition -- I know I'm right, but I can't express why. Symbol manipulation -- if the brain does not actually directly manipulate symbols, how does it learn to identify and usefully manipulate what we describe as a symbolic representation? Consciousness -- is it an emergent property at a high level, or is it like the concerted effort of a group of muscle fibers, each synapse or neuron having a little bit of consciousness? The difference between "my arm goes up" and "I raise my arm" is consciousness. Psychological development -- how does our brain and its functions develop from the overabundance of neurons that we start with to the set we grow up using? What are the incremental changes involved? Can we effectively apply this information to treatment of mental problems or to educational programs? Intelligence -- is it a measure of how many models of a certain sophistication we can usefully recognize and manipulate? Is there one basic type of intelligence or seven or more? Phantom limbs -- they are no longer there -- how does an amputee continue to have the sensations of the missing limb? Dreams -- are dreams the brain's way of retraining the nervous system while we sleep?[3] [4] Systems effects, emergent properties. How is it that unexpected abilities or functions become apparent when elements of the substrate act in concert? What structural dynamics account for these properties? The combined efforts of neural scientists and computer neural network experimenters can help eliminate inappropriate theories of mental function using the Sherlock Holmes Rule of

3. Wolf, Fred Alan. *The Dreaming Universe*. New York: Touchstone/Simon & Schuster. 1994.
4. Rahmann, *Ibid.*

Evidence: when all the impossibilities have been eliminated, whatever remains, however improbable, must be the correct answer. Neural networks can be trained to translate neural signals to movement in a prosthesis. There is a possibility of mapping visual activity in the brain to a neural network. We may be able to directly read the visual content of the brain, perhaps even recollections of things seen or merely imagined. To view the imagination at work is becoming at least a theoretical possibility. We have reached the outer limits of speculation.

Figure 02: "...the most complex object in the universe. It comprises a trillion cells, 100 billion of them neurons linked in networks that give rise to intelligence, creativity, emotion, consciousness and memory." [5]

---

[5]. Fischbach, Gerald D. "Mind and Brain". Scientific American, September 1992, 51. This article provides the bulk of the discussion and description of brain structure in this exposition.

I. A Natural Neural Network: The Brain

Composition

The brain is composed largely of neurons which are connected to each other in multiple ways, neurons which operate as tiny individual computers, summing inputs and sending out signals to other neurons, both locally and at a distance, neurons which are integrated into highly complex functional systems in a hierarchy which begins with the lowly synapse and culminates in human behavior. Analysis of these structures at a fine level is made difficult by several obstacles: (1) the system is immensely complex, composed of several hundred billion neurons; (2) the number of connections is extremely large -- as many as 10,000 terminals per neuron; (3) there are diverse modes of synaptic association between two neurons -- excitatory-inhibitory, electrical-chemical; and (4) there is a great diversity of transmitter substances, probably around 50.[6] [7] This lack of more detailed information about brain structures and their functions will make the comparison to artificial neural networks even more difficult.

Structure

The number of identifiable levels of brain structure and their significance for cognition are issues for debate. The human brain weighs three to four pounds and contains about 100 billion neurons. This immense number alone does not account for the brain's complexity:

6. Rahmann, *Ibid.* 99.
7. Fischbach. *Ibid.* 48-57. As of this writing, there are (at least) between 100 and 200.

> Although that extraordinary number is of the same order of magnitude as the number of starts in the Milky Way, it cannot account for the complexity of the brain. The liver probably contains 100 million cells, but 1,000 livers do not add up to a rich inner life.[8]

There is a great diversity of cells -- Purkinje, basket, motor neurons, etc. A Purkinje cell alone is amazingly complex. The challenge to define the function of one such cell -- or even a small group taken together -- is daunting. The brain is not one smooth net of similar cells performing identical functions.

There are some simplifications in brain structure which do facilitate analysis of structure and function: 1) Groups with similar functions are grouped together in columns (or "slabs") that extend through the thickness of the cortex. For example, a module in the visual cortex which responds to a line of a particular orientation could contain up to 100,000 cells, most of which participate in local circuits devoted to a particular function; 2) All neurons conduct information in much the same way: information travels along axons in the form of brief electrical impulses called action potentials, which measure about 100 millivolts in amplitude and one millisecond in duration. These result mainly from the exchange of potassium and sodium ions across the surface membrane of the neuron from the extracellular fluid into the cell interior or cytoplasm. At a critical potential called the threshold, an electrical charge is generated and positive feedback produces a regenerative event that forces the membrane potential to reverse sign. The sodium permeability mechanism remains refractory for a

8. Fischbach. *Ibid.* 49.

few milliseconds after each pulse which limits the rate at which action potentials can be generated to 200 per second or less. The need for signal boosting limits the maximum speed at which an impulse travels to about 100 meters per second, less than one millionth the speed at which an electrical signal moves in a copper wire. Thoughts must depend on the relative timing of impulses conducted over many axons in parallel and on the thousands of connections made by each one.

Transmitters

Communications between neurons is mediated by chemical transmitters that are released at specialized contacts called synapses, which, we will see later, are the best candidates to be the basic computational unit of the network.[9] Each neuron must continually integrate up to 10,000 synaptic inputs, which do not add up in a simple linear manner, which means each neuron is a sophisticated computer in itself.

> Transmitter receptors can be grouped into two large (and growing) superfamilies based on their amino acid sequence and on presumptions about the shape that the molecules assume as part of the cell membrane in which they are embedded . . . one receptor superfamily consists of ion channels, [which] underlie . . . changes in permeability . . . The other superfamily . . . does not form channels. Instead its members interact with a neighboring membrane protein . . . This process initiates a cascade of biochemical reactions.[10]

This chemical interaction must take place at an incredible pace in order to have a timely effect on the electrical activity in the system. The pace of chemical change must be at least that of the synaptic interactions with the neurons, faster than 100-200

9. See Appendix 4. Defining the Basic Computational Unit.
10. Fischbach. *Ibid.* 54.

transformations per second.

Plasticity

The account of structural, functional and molecular variety given so far would seem to be sufficiently complete as a basis for mental function. We have yet to consider plasticity, the tendency of synapses and neuronal circuits to change as a result of activity. "Plasticity weaves the tapestry on which the community of mental life depends."[11] The metabolic after-effects of action potentials not only encode information, they alter the circuits over which they are transmitted.

Synapse plasticity is the basis for connectionist neural models. It "multiplies the complexity provided by any fixed cast of molecular characters or cellular functions," providing an even richer substrate for mental phenomena. These changes may alter function of synapses as well as the number or location of synapses themselves. Axons generate new endings when nearby neurons cease transmission, and the terminal branches of dendritic arbors are constantly remodeled.

> Although the forces leading to plastic changes in the mature brain are frequent and ineluctable, it is important to emphasize the precision and overall stability of the wiring diagram. We could not sense the environment or move in a coordinated manner, let alone think, if it were otherwise. All studies of higher brain function must take into account the precise way in which neurons are connected to one another.[12]

11. Fischbach. *Ibid.* 54.
12. Fischbach. *Ibid.* 55.

Topographic Mapping

> Pathways in the brain have been traced by means of a variety of molecules that are transported along axons. Such reporter molecules can be visualized once the tissue is properly prepared. Connections have also been traced by fine-tipped microelectrodes positioned close enough to a nerve cell body or an axon to detect the small currents generated as an action potential passes by. Each technique has revealed ordered, topographic maps in the cerebral cortex. The body surface is represented in the postcentral gyrus of the cerebral cortex even though the cortical neurons are three synapses away from sensory receptors in the skin. Likewise, a point-to-point map of the visual world is evident in the primary visual cortex at the occipital pole at the back of the brain. Order is evident at each of the early relays on route to the cortex, and topographic order has also been found in projections from the primary cortices to higher centers.[13]

Hubel and Wiesel discovered 30 years ago that neurons in the primary visual cortex (V1) respond to line segments or edges of a particular orientation rather than to the small spots of light that activate the input neurons in the retina and lateral geniculate nucleus of the thalamus, implying that neurons in V1 are connected, via the lateral geniculate nucleus, to retinal ganglion cells that lie along a line of the preferred orientation. "There is a general topographical correspondence between a region of simulation of the sensory organ and its representation in the central nervous system . . . While maintaining topographical coherency, the nature of the information stored in the brain can change as it progresses from one way-station to the next."[14]

13. Fischbach. *Ibid.* 55.
14. Cook, Norman D. *The Brain Code.* New York: Methuen, 1986, 33.

Brain Code

We know the anatomy of the major sensory and motor systems in some detail. However, the pattern of connections within the intervening association cortices and the large subcortical nuclei of the cerebral hemispheres is not clearly defined. Visual information is mapped so that the retinal image in the eye maintains its basic configuration as the pattern of stimulation is transferred from the retina to the lateral geniculate body in the thalamus to striate cortex. Somatosensory information is also topologically intact, with the entire surface of the skin being represented contiguously and specially detailed representation for the hand. Cook makes a distinction here between study of the neuron and its synaptic activity, which he terms *neurophysiology*, and the study of larger aggregates of neural activity which account for specific cognitive functions or for functions which taken together account for mental and physical behavior, which he terms the *Brain Code* or *neuropsychology*.[15]

The pattern of information flow in the brain during the performance of mental tasks cannot easily be determined by anatomic studies of the circuit diagram or by studies of plasticity. Neural correlates of higher mental functions are being sought directly in awake primates trained to perform tasks that require judgment, planning or memory, or all three capacities.

---

15. Cook. *Ibid.* xiii.

Hierarchies

> One of the most important principles is that sensory systems are arranged in a hierarchical manner. That is, neurons respond to increasingly abstract aspects of complex stimuli as the distance -- measured in numbers of synapses from the source -- grows. The fact that neurons in V1 respond to lines rather than spots makes the case. Another important principle . . . is that information does not travel along a single pathway. Rather, different features of a single percept are processed in parallel pathways . . . the movement, color and shape of a tennis ball are processed in different cortical visual centers.[16]

The auditory system has a similar topographical organization, with nearby structures responding to similar frequencies. But, in the barn owl, phase and amplitude signals which account respectively for location along the azimuth and elevation are processed in different pathways through three synaptic relays in the brain. It seems likely that this type of parallel processing characterizes other sensory systems, association cortices and motor pathways as well. For years, psychologists have seen the normal brain as a black box whose inner workings were the speculations of many different schools of thought. Occasionally some area of a brain would be damaged, and the corresponding loss of function would give a clue as to the function performed by that part of the brain. In this way Broca's Area and Wernicke's Area were respectively found to process speech production and language comprehension.[17] We are now reaching the point at which it is becoming possible to reach deeper into the structure and to more clearly identify the functions of smaller and smaller local

---

16. Fischbach. *Ibid.* 56.
17. Cook. *Ibid.* 19.

structures in the brain.

Cerebral Hemispheres

> Where is the information reassembled? When does the subject become aware of the approaching ball? The receptive fields of neurons in higher centers are larger than those found in earlier relay stations, so they monitor a larger fraction of the external world. Zeki describes a model that depends on feedback connections from cells with large receptive fields to the cells in the primary visual cortex that have high spatial resolution. Such feedback circuits might coordinate the activity of cells in the primary cortex that have high spatial resolution and cells that respond to more abstract features of the stimulus no matter where it is located. Francis Crick and Christof Koch address the role in visual awareness of a 40-cycle-per-second oscillation in firing rate that is observed throughout the cortex . . . [which] may synchronize the firing of neurons that respond to different components of a perceptual scene and hence may be a direct neural correlate of awareness.[18]

How do we close the big loop?[19] I mean the feedback loop that accounts for hand-eye coordination, that doesn't just see the ball coming, but takes steps to intercept it. Fischbach asks where it all comes together. Cook has a suggestion for where to look: the corpus callosum. The brain is composed of a number of identifiable parts, the most identifiable being the two lateral hemispheres. The cerebral cortex is outer 2-3 mm of the entire surface of the cerebral hemispheres. Its surface forms a series of elevated ridges, or gyri, separated by shallow depression, called sulci, or deeper grooves, called fissures.[20] The total surface area is approximately 25 square

18. Fischbach. *Ibid.* 56.
19. Dreyfus. Hubert L. "Phenomenology of Embodiment." *The Electronic Journal of Analytic Philosophy*, 4 (Spring 1996). Bloomington: Indiana University. 1996. ¶ 32.
20. Martini, Frederick H. *Fundamentals of Anatomy and Physiology, 3d ed.* Englewood Cliffs: Prentice Hall. 1995, 468.

feet of flat surface -- it is significant that the topology of the cortex is actually two- rather than three-dimensional.

The prefrontal cortex is involved with conscious intellectual functions. The frontal lobe includes the primary motor cortex, and processes voluntary control of skeletal muscles. The parietal lobe is associated with conscious perception of touch, pressure, vibration, pain, temperature, and taste. The occipital lobe contains the primary visual cortex, and the temporal lobe accounts for auditory and olfactory cortices. The sensory and motor regions of the cortex are connected to nearby association areas that interpret incoming data or coordinate motor response. Integrative centers include the prefrontal cortex, which integrates information from sensory association areas and performs abstract intellectual functions, such as predicting the consequences of possible responses, the general interpretive area (for language and mathematical calculation), which is usually confined to one (the left) hemisphere, and the speech center.[21]

The corpus callosum contains over 200 million axons, carrying an estimated 4 billion impulses per second.[22] Its central location and function as the carrier of signals between hemispheres are what made it the focus of Cook's *Brain Code* book. He lists the accompanying table (Table 1.) of *dichotomies attributed to the cerebral hemispheres*:[23] Martini calls the left hemisphere the *categorical hemisphere*, the right

21. Martini. *Ibid.* 470-2.
22. Martini. *Ibid.* 469.
23. Cook. *Ibid.* 18.

hemisphere the *representational hemisphere*.[24]

Cook argues the necessity of one hemisphere to dominate in certain activities -- for example, the coordination of movement and expression. You can't clap your hands very well if each hand is controlled independently, so one side must coordinate the effort. The communication for this coordination goes through the corpus callosum.

There is still much to be found out about brain function in the various areas. A recent study challenged the traditional theory of cerebellar function, showing that the cerebellum is not directly responsible for fine motor control, but is involved "in sensory discrimination rather than in movement per se."[25] And we still wonder where it all comes together. Our most persistent interest in this pursuit of knowledge is the source of consciousness, which we frequently associate with cognition. But the brain performs many unconscious functions at many levels. One wonders just how many of those levels we are privileged to experience as consciousness. Are we limited to one? Do we occasionally get a glimpse into others, as when we are between dream and wakefulness?[26]

Firing Rate and Vectors

How many neurons must change their firing rate to signal a

24. Martini. *Ibid.* 471-2.
25. Gao, Jia-Hong, Lawrence M. Parsons, James M. Bower, Jinhu Xiong, Jinqi Li, peter T. Fox. "Cerebellum Implicated in Sensory Acquisition and Discrimination Rather than Motor Control." Science VOL N.d. 1996, 1-3.
26. Wolf. *Ibid.*

Table 01. Dichotomies attributed to the cerebral hemispheres

Left Hemisphere	*Right Hemisphere*
Experimentally Derived	
sequential	parallel/simultaneous
analytic	synthetic
linguistic	visuospatial
verbal	nonverbal
focal	global
linear	spatial
mathematical	geometric
Philosophically Motivated	
expressive	perceptive
propositional	appositional
symbolic	imaginative
intellection	emotion
reason	affect
discrete	holistic
active	passive

coherent percept or gestalt? The most extreme view holds that one cell may do the job. Is there one face cell per face? Such a supposition seems unlikely on first principles: we lose thousands of neurons every day, so overcommitment to one would be unwise. A more compelling argument comes from recent experiments that have shown face cells to be broadly tuned, responding to faces with similar features rather than to one face alone. The number of neurons that must be activated before recognition emerges is not known, but the data are consistent with a sparse coding rather than global or diffuse activation.

Face cells have their counterparts on the motor side. "Command" neurons have been identified in certain invertebrates that trigger all-or-none, fixed-action patterns, such as stereotypical escape behaviors . . . like face cells in the temporal lobe, individual motor cortex neurons are broadly tuned.

The vector obtained by summing the firing frequencies of many neurons is better correlated with the direction of movement than is the activity of any individual cell. The vector becomes evident several milliseconds *before* the appropriate muscles contract and the arm actually moves. It must be a sign of motor planning. The vector is usually derived from less than 100 neurons, so sparse coding may be the rule in the motor cortex as it is in the temporal sulcus.[27]

Experimental Evidence

Studies are underway to produce mental phenomena by focal electrical stimulation. This harkens back to the memory stimulation exercises by Wilder Penfield, in which he evoked vivid memories from patients by stimulating their brains directly electronically while they were undergoing brain surgery.[28] Strokes and other unfortunate "experiments of nature" have also provided important insights regarding neural correlates of mental phenomena.

The future of cognitive neuroscience depends on our ability to study the living human brain. PET and MRI hold great promise in this regard [29]. The brain is never completely at rest. At present, neither technique provides the spatial resolution to visualize single cortical columns.

Fischbach believes we can expect advances at an increasing rate on all levels of investigation relevant to the mind: We will soon know exactly how many transmitters and transmitter receptors there are in the brain and where each one is concentrated . We will also have a more complete picture of neurotransmitter actions, including multiple interactions of jointly released modulators. And we will learn much more about molecules that affect neuronal differentiation and degeneration. The great

27. Fischbach. *Ibid.* 57.
28. Penfield, Wilder. *The Mystery of the Mind.* Princeton: Princeton University Press, 1975.
29. Gao, et al. *Ibid.* 1-3.

challenge is to determine how these molecules modulate the functional wiring diagram of the brain and how this functional nerve network gives rise to mental phenomena.

Emergence

Ultimately, it will be essential to specify what exactly it means to say that mental events are correlated with electrical signals. Is the mind an emergent property of the brain's electrical and metabolic activity? An emergent property is one that cannot be accounted for solely by considering the component parts one at a time . . . biological explanations of mental events may become evident once the component neural functions are more clearly defined. We will then have a more appropriate vocabulary for describing the emergent mind.[30] *Emergence* is frequently discussed when systems produce behavior which can not be expected from a minute analysis of all of the subcomponents, in other words, when a system exhibits behavior which can not be predicted by examining the properties of its constituent parts. Consider *sand*. Consider *dune*. Even if you know all properties of each individual grain, and the dynamic relations among the grains, and properties of the grains when aggregated, you do not have enough information at that level to explain an individual dune and its systemic properties.[31] An appropriate explanation can be made at a higher level of organization, one that recognizes the applicable models for friction and fluid flow of granular aggregations. Once all factors have been considered and weighed,

30. Fischbach. *Ibid.* 57.
31. Nagel, Ernest. *The Structure of Science*. New York: Harcourt, Brace & World, Inc. 1961. 366-80.

both at the substrate level and at the local level, environmental factors may have a role which is yet unexplained. When we understand all the factors affecting a system, we hope we will have a set of laws which completely predict its behavior, but until such a comprehensive system description is confirmed, we may profitably consider unexplained system phenomena as *emergent*. Supposedly, equally mysterious properties *emerge* from other types of aggregates, such as neuron bundles. Consider muscle. Consider muscle fiber. Here, the properties of the muscle inhere to each individual fiber, and the power of the muscle is merely an aggregation of smaller powers.[32] Is the mind merely an aggregate of mindlets? Or is there something that must be explained by appeal to a higher level model? We are ever closer to plausible explanations for the physical basis of mental phenomena, but much more experimental evidence must be processed before we can claim to know how these systems work. In the meantime, the interplay between investigations of artificial and natural neural networks provides insights for all neural networks.

There is a brilliant visualization of what one layer of a neural network might look like and how it might function in "Pinscreen Excerpts," featuring the animation work of Alexander Alexeieff.[33] A pinscreen consists of thousands of dark pins closely clustered and thrust through a white flat background. The height of the pins above the

---

32. Shepherd, Gordon M. "The Significance of Real Neuron Architecture for Neural Network Simulations." Chapter 8 in Eric Schwartz, *Computational Neuroscience*, (q.v.)
33. Alexander, R. McNeill. *Exploring Biomechanics: Animals in Motion*. New York: Scientific American Library. 1992.

surface and the angle of lighting determines how much white shows through, creating a very finely-grained black-and-white picture. The pins are adjusted frame after frame to create a moving picture in the same way that other animation techniques create an impression of movement through gradual change. The difference is that the subtleties of the representation are made using small changes in individual elements which are not simply on or off or in a certain color or contrast, but which are arbitrarily finely adjustable, and which yet form a readily identifiable whole. To an outside observer, of course, this representation is no more enlightening than a table of node weights. But some mechanism monitors these weights, compares the results of applying the current network, and retrains the network to respond with a different weighting on the next input. A picture emerges from a pinscreen, but this is due to conscious human activity. We are not conscious of the activity which governs our own natural neural network. We trust that, at some level, there is an explanation for everything. As advances are made, not just in neurological science, but in other sciences, both hard and soft, for clues which will lead us to a clear understanding of what is really meant when we say we think, or that we have thoughts, or that we perceive, or understand. New ground is being broken in neural studies which show us that the brain, even now, is not necessarily doing exactly what we think it is.[34] Studies in physics show us that there is a natural tendency towards self-organization in dynamic systems.[35] Sand.

---

34. Gao, et al. *Ibid.*
35. Umbanhowar, Paul B., Francisco Melo & Harry L. Swinney. "Localized excitations in a vertically vibrated granular layer." *Nature* 382, 793-796. 1996. I suspect there may be similar activities in natural neural

Dune. Two different models constituted of the same substance, the one the sole substrate of the other. Pile up the sand -- which particle made the dune? How we see it and understand it as a model differs as our model of sand and our model of dune differ. Our model of dune has nothing directly to do with our model of sand particle. A dune has topological implications completely missing in the sand model.

A recent issue of Scientific American[36] showed magnified images of sand particles form many locations around the world. It showed the particles to be extremely varied, and of many beautiful designs. Our normal sight can barely differentiate one particle from another. This beauty of form is missing from our model of sand.

Anything seen macroscopically presents a different model than that seen microscopically, with the exception of fractal patterns (little fleas have lesser fleas . . .) The significance of a model at one level need not have any bearing on the significance of a higher level model, even though it is the sole substrate of that model.[37] On the other hand, some models are so close together that operations on one readily translate into operations on the other. Many such models are mathematically expressible.[38] Tools are available to help us study actual natural neural networks in action: CAT scan, PET, MRI, etc. The use of artificial neural networks as not simply beneficiaries

network to those in the granular layers investigated here.

36. Mack, Walter N. and Elizabeth A. Leistikow. "Sands of the World." *Scientific American*. Volume 275, Number 2, August 1996. 62-67.

37. See Appendix 4. Defining the Basic Computational Unit.

38. See Appendix 2. Emergent Properties.

of connectionist theory, but also as tools of experimentation can enhance our knowledge of neural network behavior economically and safely.

II - Artificial Neural Networks

Introduction

The introduction of the idea of neurons as structural constituents of the brain is attributed to Ramón y Cajál (1911).[39] Eighty-five years later, many questions remain concerning just how the system performs the work it performs. Against a background of controversy over just what level of operation is appropriate for discussion of the various levels of function of the natural neural network of the brain, computer scientists are trying to emulate the brain's functions by simulating its operation in artificial neural networks either in a massively parallel computing machine or in a von Neumann architecture which can simulate parallel processing. Neurons are much slower than silicon logic gates, five to six orders of magnitude slower. Events in a silicon chip cycle in the nanosecond (10^{-9} s) range, whereas neural events are in the millisecond (10^{-3} s) range. The brain makes up for the individual neuron's slow rate of operation by having on the order of 10 billion neurons in the human cortex, with 60 trillion synapses or connections among them.[40] The net result is that the brain is an enormously efficient structure, with an *energetic efficiency* of approximately 10^{-16}

39 Haykin. *Ibid.* 1. This source represents the bulk of substantive information on the structure and theory of artificial neural networks in Part II of this report.

40. Churchland, Paul M. "A Deeper Unity: Some Feyerabendian Themes in Neurocomputational Form." Chapter 4 in Steven Davis (q.v.). *Connectionism: Theory and Practice.* 41.

joules (J) per operation per second, whereas the corresponding value for the best computers currently (late 1996) in use is about 10^{-6} joules per operation per second. It is possible on current computers to simulate only a small portion of this complexity in any given instance. It is not yet possible to copy the exact structure of natural networks in artificial networks. The brain is a highly *complex, nonlinear, and parallel computer* (information-processing system) which performs complex perceptual recognition tasks in 100-200 ms, whereas tasks of much lesser complexity will take days on a huge conventional computer. We are currently limited in the extent to which we may attempt to emulate the brain by the availability of, both in terms of computational structure and of raw memory.

At birth the brain has many more neurons, and great flexibility in how they are to be connected. Experience -- perception and interaction with the real world -- leads to restructuring of the brain's circuits, with many neurons disappearing while new synapses continue to be formed. Although development continues well beyond that stage, the most dramatic development (i.e., hard-wiring) of the human brain takes place in the first two years. During this early stage of development, about 1 million synapses are formed per second. The synapse is arguably the *basic computational unit* of the brain. Rahmann states " . . . a neuron, as the elemental unit of the nervous system, constitutes both an ontogenetic (developmental) and a physical (i.e., trophic and functional) unit that is responsible for perception, processing, transmission, and, above all, storage of information in an organism." [41]

Synapses are elementary structural and functional units that mediate the interactions between neurons. The most common kind of synapse is a *chemical synapse*, which operates as follows. A presynaptic process liberates a transmitter substance that diffuses across the synaptic junction between neurons and then acts on a postsynaptic process. Thus a synapse converts a presynaptic electrical signal into a chemical signal and then back into a postsynaptic electrical signal. In electrical terminology, such an element is said to be a *nonreciprocal two-port device*. In traditional descriptions of neural organization, it is assumed that a synapse is a simple connection that can impose *excitation* or *inhibition*, but not both on the receptive neuron.[42]

Haykin identifies a developing neuron with a plastic brain:

> "*Plasticity* permits the developing nervous system to adapt to its surrounding environment. In an adult brain, plasticity may be accounted for by two mechanisms: the creation of new synaptic connections between neurons, and the modification of existing synapses. *Axons*, the transmission lines, and *dendrites*, the receptive zones, constitute two types of cell filaments that are distinguished on morphological grounds; and axon has a smoother surface, fewer branches, and greater length, whereas a dendrite (so called because of its resemblance to a tree) has an irregular surface and more branches. Neurons come in a variety of shapes and sizes in different parts of the brain . . . In its most general form, a *neural network* is a machine that is designed to model the way in which the brain performs a particular task or function of interest; the network is usually implemented using electronic components or simulated in software on a digital

41. Rahmann. *Ibid.* 2.
42. Haykin. *Ibid.* 2.

computer . . . a neural network viewed as an adaptive machine:
A neural network is a massively parallel distributed processor that has a natural propensity for storing experiential knowledge and making it available for use. It resembles the brain in two respects:
1. Knowledge is acquired by the network through a learning process.
2. Interneuron connection strengths known as synaptic weights are used to store the knowledge. (adapted from Aleksander and Morton (1990))[43]

The procedure used to perform the learning process is called a *learning algorithm*, the function of which is to modify the synaptic weights of the network in an orderly fashion so as to attain a desired design objective.

Neural networks are also referred to in the literature as neurocomputers, connectionist networks, or parallel distributed processors. A number of benefits are claimed for neural networks. Among them, Haykin cites the following nine:[44]

1. *Nonlinearity* -- In actually using a parallel processing mechanism or simulating one on a sequential computer, it is possible to mimic natural neural networks. Just how valid this emulation may be is a crucial part of our later discussion.

2. *Input-Output Mapping* -- By using *supervised learning*, the network learns from examples, repeating *training examples* until a steady state is reached, and then applying the network to *task examples*, theoretically emulating natural neural network functions.

3. *Adaptivity* -- Networks can be retrained to perform new tasks or modify

43. Haykin. *Ibid.* 2.
44. Haykin. *Ibid.* 4-6.

their performance of old tasks. This flexibility is desirable in many contexts, but adaptivity is not always consistent with robustness. Too great a sensitivity to random noise or other spurious interference can degrade performance unnecessarily. This is the *stability--plasticity dilemma.*

4. *Evidential Response* -- In addition to which selection to make, the system can determine the level of confidence in that particular choice.

5. *Contextual Information* -- is dealt with naturally in neural networks, since knowledge is represented in the very structure and activation state of the network.

6. *Fault Tolerance* -- Neural networks provide a *graceful degradation* in performance when the network is damaged, and may continue to give reliable responses long after a symbolic program will have failed altogether.

7. *VLSI Implementability* -- Implementation using *Very Large Scale Integrated* technology makes possible real-time applications involving pattern recognition, signal processing, and control.

8. *Uniformity of Analysis and design* -- The *universality* of neural networks as information processors facilitates the use of the neuron model as an ingredient common to all neural networks, makes it possible to share theories and learning algorithms, and allows construction of modular networks through a seamless integration of modules.

9. *Neurobiological Analogy* -- Network design is "motivated by analogy with

the brain, which is a living proof that fault-tolerant parallel processing is not only physically possible but also fast and powerful. Neurobiologists look to (artificial) neural networks as a research tool for the interpretation of neurobiological phenomena. For example, neural networks have been used to provide insight on the development of premotor circuits in the oculomotor system (responsible for eye movements) and the manner in which they process signals. On the other hand, engineers look to neurobiology for new ideas to solve problems more complex than those based on conventional hard-wired design techniques . . ."

In Part I, I identified an eight-level structural organization for natural neural networks. Haykin identifies two types of structural model which may be applied, a three-level model, and an eight-level model. The three-level model is comprised of an *input* level, the *stimulus*, a *processing* level which includes a *neural network* sandwiched between *receptors* and *effectors*, and a *response*:

Stimulus $\rightarrow$ [Receptors $\leftrightarrow$ Neural net $\leftrightarrow$ Effectors] $\rightarrow$ Response.

This model seems appropriate for psychological applications among others. This model includes feedback between the receptors and the network, and between effectors and the network. The extent and type of feedback that is appropriate in neural networks is a matter for debate. Haykin's eight-level analysis is very similar to that of the natural network scientists:[45]

45. Haykin. *Ibid.* 6-7.

1. Central nervous system

2. Interregional circuits

3. Local circuits

4. Neurons

5. Dendritic trees

6. Neural microcircuits

7. Synapses

8. Molecules

A neural microcircuit refers to an assembly of synapses organized into patterns of connectivity so as to produce a functional operation of interest. These are grouped to form *dendritic subunits* within the *dendritic trees* of individual neurons. The whole *neuron* contains several dendritic subunits. At the next level of complexity are *local circuits* made up of neurons with similar or different properties which perform operations characteristic of a localized region in the brain. Then come *interregional circuits* made up of pathways, columns, and topographic maps, which involve multiple regions located in different parts of the brain. *Topographic maps* are organized to respond to incoming sensory information. Whether this topography is an essential part of cognition is also a matter for contention. These maps are often arranged in sheets, as in the superior colliculus, where the visual, auditory, and somatosensory maps are stacked in adjacent layers in such a way that stimuli from corresponding points in space lie above each other. Finally, the topographic maps, and other interregional circuits mediate specific types of behavior in the *central nervous system (CNS)*.

It is important to recognize that the structural levels of organization which I described are a *unique* characteristic of the brain. They are not found in a digital computer, and we are far from realizing them with artificial neural networks. However, we are gradually making our way toward a hierarchy of computational levels similar to that just described. The artificial neurons we use to build our neural networks are truly primitive in comparison to those found in the human brain. We are not presently able to design a neural network with anywhere near the complexity of the local circuits and the interregional circuits in the brain.

Neuron Models

For artificial neural networks, a *neuron* is a fundamental information-processing unit with three basic elements:

1. A set of *synapses* or *connecting links*, each of which is characterized by a weight or strength of its own. The weight is positive if the associated synapse is excitatory; it is negative if the synapse is inhibitory.

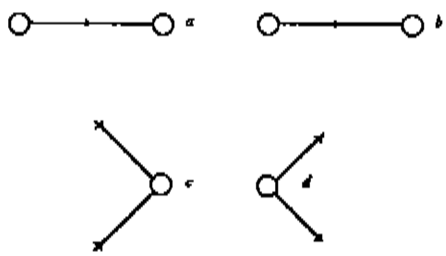

Figure 03. Basic rules for constructing signal-flow graphs.
(Based on Haykin, p.14)

2. An *adder* for summing the input signals, weighted by the respective synapses of the neuron; the operations described here

constitute a *linear combiner.*

3. An *activation function* for limiting the amplitude of the output of a neuron. The activation function is also referred to in the literature as a *squashing function*. Activation functions vary from straightforward threshold functions which have a sharp effective signal cut-off, to sigmoid functions which can have a gradual effect or a sharp effect depending on the slope parameter. The particular mathematical functions used to calculate activation are of little philosophical import except at the sub-synapse level. As we saw in the description of brain function, the synapse is the basic computational unit, and the substrate of the basic computational unit may be safely ignored for the purposes of this paper. For this reason we will eschew a study of activation function mathematics and focus more on structural detail and activation patterns.

Haykin[46] suggests simplifying the appearance of the model of an artificial neuron by using the idea of signal -flow graphs with a well-defined set of rules (Fig. 03). A *signal-flow graph* is a network of directed links (branches) that are interconnected at certain points called nodes. The flow of signals is dictated by three basic rules:

RULE 1. Signal flows along a link only in one direction defined by the arrow on the link. Two different types of links may be distinguished: (a) *Synaptic links,*

46. Haykin. *Ibid.* 13.

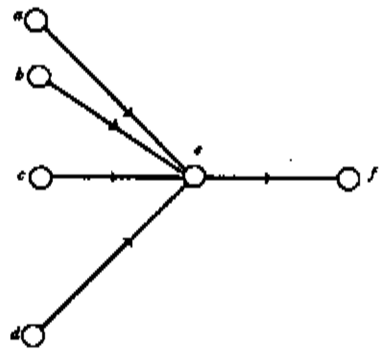

Figure 04: Typical neuron.

governed by a *linear* input-output relation, and (b) *Activation links*, governed in general by a *nonlinear* input-output relation.

RULE 2. A node signal equals the algebraic sum of all signals entering the pertinent node via the incoming links. This rule (c) represents *synaptic convergence* or *fan-in*.

RULE 3. The signal at a node is transmitted to each outgoing link originating from that node, with the transmission being entirely independent of the transfer functions of the outgoing links. This rule (d) represents *synaptic divergence* or *fan-out*.

This signal-flow graph model (Fig. 04) results in the following mathematical definition of a neural network:[47]

> *A neural network is a directed graph consisting of nodes with interconnecting synaptic and activation links, and which is characterized by four properties:*
> 1. *Each neuron is represented by a set of linear synaptic links, an externally applied threshold, and a nonlinear activation link. The threshold is represented by a synaptic link with an input signal fixed at a value of -1.*
> 2. *The synaptic links of a neuron weight their respective input signals.*

47. Haykin. *Ibid.* 14-15.

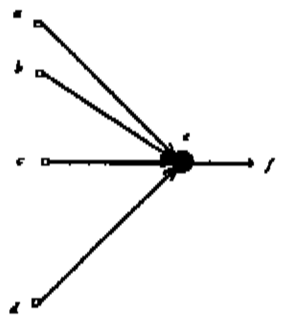

Figure 05. Architectural Graph of a Neuron

3. *The weighted sum of the input signals defines the total internal activity level of the neuron in question.*
4. *The activation link squashes the internal activity level of the neuron to produce an output that represents the state variable of the neuron.*

Such a directed graph is defined as *complete* in that it describes not only the signal flow from neuron to neuron, but also the signal flow inside each neuron. Other representations are also used, such as an *architectural graph* (Fig. 05), which is the result of omitting the details of signal flow inside the individual neurons. Such a directed graph is said to be partially complete, and is characterized as follows:
1. Source nodes supply input signals to the graph.
2. Each neuron is represented by a single node called a *computation node*.

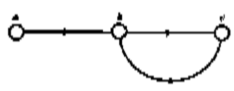

Figure 06. Single-loop feedback

3. The *communication links* interconnecting the source and computation nodes of the graph carry no weight; they merely provided the directions of signal flow in the graph.[48]

Feedback.

Feedback exists in a dynamic system when the output of an element in the system influences the input applied to that particular element, giving rise to one or more closed paths for the transmission of signals around the system. It occurs in almost every part of the nervous system of every animal, and it plays a major role in the study of the special class of neural networks known as *recurrent networks*.[49] There are various mathematical models both implementing and for analyzing feedback. The model selected may result in a stable system in which the signal is convergent, or in an unstable, divergent system. The former has the attribute of *infinite memory*, in the sense that the output of the system depends on samples

48. Haykin. *Ibid.* 15.
49. Haykin. *Ibid.* 15.

of the input extending into the infinite past. Moreover, the memory is fading in that the influence of a past sample is reduced exponentially with time *n*.[50]

Network Architectures[51]

The structure of a neural network is intimately linked with the learning algorithm used to train the network. We may therefore speak of learning algorithms (rules) used in the design of neural networks as being structured.

We identify four different classes of network architectures:

1. Single-layer Feedforward Networks

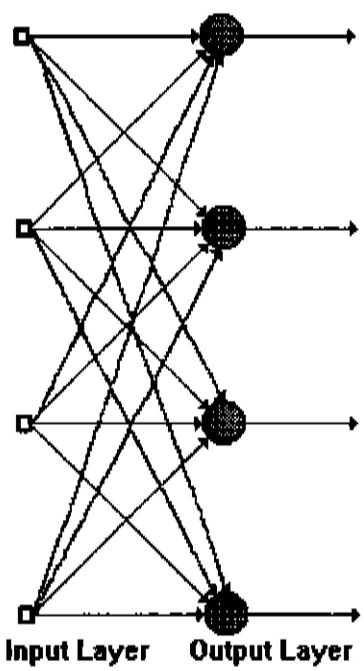

Figure 07. Single-layer Feedforward Network

A *layered* neural network is a network of neurons organized in the form of layers. In the simplest form of a layered network, we just have an *input layer* of source nodes that projects onto an *output layer* of neurons (computational nodes), but not vice versa. In other words, this network is strictly of a *feedforward* type. Such a network is called a *single-layer network* (Fig. 07), with the designation "single layer" referring to the output layer of

50. Haykin. *Ibid.* 17.
48. Haykin. *Ibid.* 18.

computational nodes (neurons). In other words, we do not count the input layer of source nodes, because no computation is performed there.

A linear associative memory is an example of a single-layer neural network. In such an application, the network associates an output pattern (vector) with an input pattern (vector), and any information is stored in the network by virtue of modifications made to the synaptic weights of the network.

2. Multilayered Feedforward Networks

The second class of a feedforward neural network distinguishes itself by the presence of one or more *hidden layers*, whose computational nodes are correspondingly called hidden neurons or hidden units (Fig. 08). The function of the hidden neurons is to intervene between the external input and the network output. Hidden layers allow a system to create its *own* internal representation. By adding one or more hidden layers, the network is enabled to extract higher-order statistics,

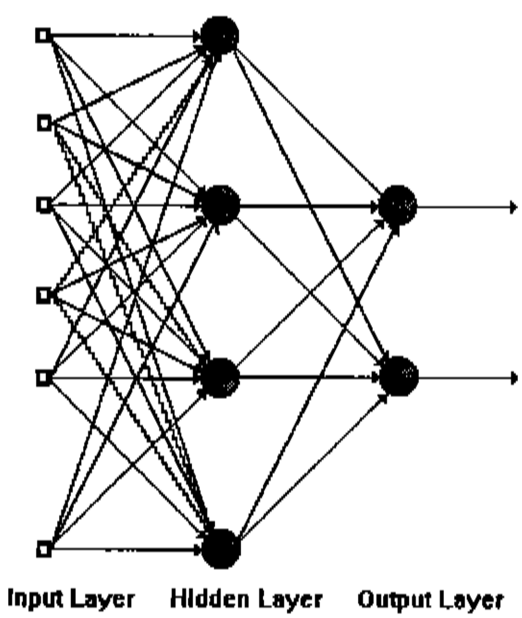

Figure 08. Feedforward Network with Hidden Layer

for (in a rather loose sense) the network acquires a global perspective despite its local connectivity by virtue of the extra set of synaptic connections and the extra dimension of neural interactions.[52] The ability of hidden neurons to extract higher-order statistics is particularly valuable when the size of the input layer is large.

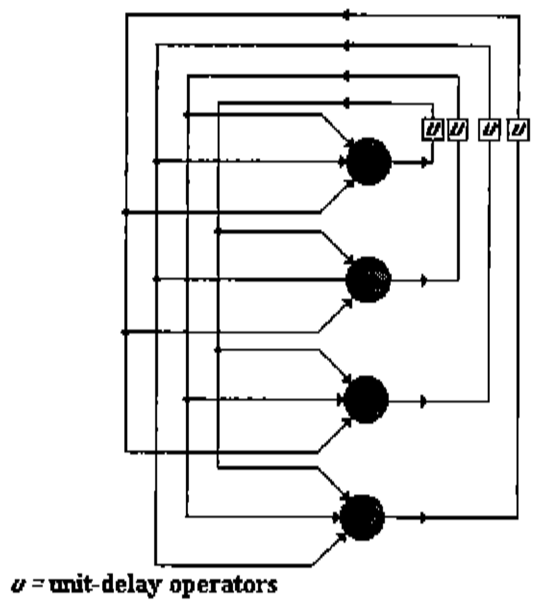

Figure 09. Recurrent network sans self-feedback.

3. Recurrent Networks

A *recurrent neural network* has at least one *feedback* loop (Fig. 09). It may consist of no more than a single layer of neurons, with each neuron feeding its output signal back to the inputs of all the other neurons. This form may contain no self-feedback loops, and may have no hidden layers. A recurrent network with hidden neurons would combine the graphical features of figures 06 and 07, adding a self-feedback connection.

4. Lattice Structures

A *lattice* consists of a one-dimensional, two-dimensional, or higher-

52. Haykin, *Ibid.* 19.

dimensional array of neurons with a corresponding set of source nodes that supply the input signals to the array; the dimension of the lattice refers to the number of the dimensions of the space in which the graph lies. A graphical representation of this structure would layer features found in figure 07.

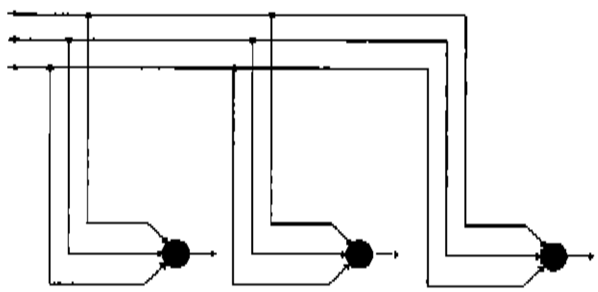

Figure 10. A one-dimensional 3-neuron lattice

Even the most elaborate of these structures is merely a simplified instantiation of a simplification of a subsection of brain structure. If these networks are viewed functionally and compared to natural networks, the question remains whether we have completely accounted for the functionality of the natural network to which we make comparison. We are still severely limited in the number of neurons and structures we are able to simulate. We are not able to completely and accurately model the activity of neurotransmitters. We are not able to accurately represent the thousands of possible connections for each neuron in even a very small network. What we lack in detail capability we hope to make up for in demonstration of principle. The fact is that neural networks -- both natural and artificial -- are able to process input in a meaningful way which is different from that of a symbol-driven mechanism. Our challenge is to determine if the artificial network is a suitable model for study of the natural network.

Knowledge Representation

Notions of intelligence and knowledge are the meat of epistemology. Modern theorists are defining intelligence in multiple ways. In addition to the traditional standards of intellectual performance, such things as kinesthetic and social intelligence have been identified.[53] Intelligence as it is normally considered by philosophy includes an evaluation of one's ability to identify and process propositional truth.[54] My own definition of intelligence includes the ability to recognize and effectively manipulate models. The intellectual abilities of a neural network are limited to its ability to differentiate one input from another. An artificial neural networks intellect is purely behavioral. There are no internal symbolic rules beyond those implicit in the structure of the network. The weights of the nodes are the basis for system behavior: the training algorithm governs the ability of the network to learn positive and negative examples of whatever paradigm is being presented. The presentation of appropriate paradigms and negative examples governs whether the network learns appropriate models. Any discussion of intellect which involves consideration and application of true propositions implicitly includes the notion of doing so *consciously*. Artificial neural networks have been considered by various theorists, but no completely

53. Biehl, Soni. "Multiple Intelligences." In Manuscript.
54. Ryle, Gilbert. *The Concept of Mind*. Chicago: The University of Chicago Press. 1949. 25-32.

satisfactory demonstrable theory has developed yet.[55] [56] If there is a ghost in this machine, it is not showing itself.[57] The assumption of consciousness in discussions of knowledge is problematic: if we cannot say that an artificial neural network provides a conscious context for its operations, then how can we say that it "knows" anything in the traditional sense of "knowing how" or "knowing that"?[58] In the sense that it has been trained to respond appropriately, an artificial neural network may be said to "know how"; in the sense that its inputs have been realistic in representing the real-world environment, the same network may be said to "know that" something is the case, or that some fact is correct as stated. The measure of these abilities is behavioral: we train a network, then we test it. Its behavior tells us how well we have done in prescribing the training algorithm, and how well we have chosen our training paradigms. This is true of a neural network: *it behaves as if it believes* its training inputs are true. It doesn't *actually* believe or know anything in the sense of being conscious or making informed decisions: a neural network provides a structural representation which exhibits certain behaviors which appear to a conscious being to be intellective. Its knowledge is not only *in* its structure, it *is* its structure. My own definition of knowledge is the retained and available collection of models one has

55. Crick, Francis. "Visual perception: rivalry and consciousness." *Nature*, Vol 379, 8 February 1996. 485-6. Crick, Francis and Cristof Koch. "The Problem of Consciousness." *Scientific American*, Vol 267, Nr. 3, September 1992. 152-159.

---. "Towards a neurobiological theory of consciousness." *The Neurosciences*, Vol 2, 1990: 263-275.

56. Revonsuo, Antti, Matti Kamppinen and Seppo Sajama. *Consciousness in Philosophy and Cognitive Neuroscience*. Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers. 1994.

57. Ryle. *Ibid.* 32.

58. Ryle. *Ibid.* 32.

recognized. In my own definition intellect is knowledge in action, the successful application of models. So we speak of knowledge and knowledge representation in artificial neural networks in a sense not usual to philosophy. Haykin repeats Fischler and Firschein's definition: knowledge refers to stored information or models used by a person or machine to interpret, predict, and appropriately respond to the outside world.[59] The two primary characteristics of *knowledge representation* are: (1) what information is made explicit; and (2) how the information is physically encoded for subsequent use. Knowledge representation is goal directed. In real-world applications, a good solution depends on a good representation of knowledge. Frequently, the possible forms of representation from the inputs to internal network parameters are quite diverse, making the development of a satisfactory solution through neural networks a serious design challenge.

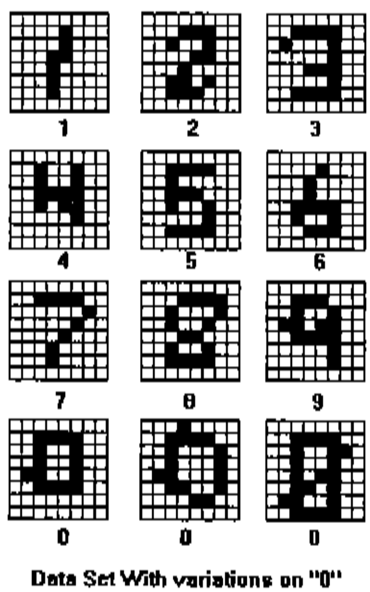

Figure 11. Digit set with sample variations

A neural network task is to learn a model of the environment in which it is embedded, and to maintain the model with

59. Haykin. *Ibid.* 22.

sufficient consistency with the real world so that it achieves the goals of the application of interest. Haykin includes two types of information under the rubric "knowledge of the world":

> 1. The known world state, represented by facts about what is and what has been known; this form of knowledge is referred to as *prior information*.
> 2. Observations (measurements) of the world, obtained by means of sensors designed to probe the environment in which the neural network is supposed to operate. Ordinarily, these observations are inherently noisy, being subject to errors due to sensor noise and system imperfections. In any event, the observations so obtained provide the pool of information from which the examples used to train the neural network are drawn.
> Each *example* consists of an input-output pair: an *input signal* and the corresponding *desired response* for the neural network. Thus, a set of examples represents knowledge about the environment of interest.
> Given a set of examples [such as hand-written digits], the design of a neural network may proceed as follows:
> First, an appropriate architecture is selected for the neural network, with an input layer consisting of source nodes equal in number to the pixels of an input image (Fig. 10), and an output layer consisting of 10 neurons (one for each digit). A subset of examples is then used to train the network by means of a suitable algorithm. This phase of the network design is called *learning*.
> Second, the recognition performance of the trained network is tested with data that has never been seen before. Specifically, an input image is presented to the network, but this time it is not told the identity of the digit to which that particular image belongs. The performance of the network is then assessed by comparing the digit recognition reported by the network with the actual identity of the digit in question. This second phase of the network operation is called *generalization*, a term borrowed from psychology.[60]

This is significantly and fundamentally different from the design of a classical

60. Haykin. *Ibid.* 23.

information-processing pattern classifier which would perform a similar function: in that design we usually proceed by first formulating a mathematical model of environmental observations, validating the models with real data, and then building the design on the basis of the model. The design of a neural network is based directly on real data, with the *data set being permitted to speak for itself*, not only providing an implicit model of the environment in which it is embedded, but also performing the information-processing function of interest.

The examples used to train a neural network may include both *positive* and *negative* examples. In a neural network, knowledge representation of the surrounding environment is defined by the values taken on by its synaptic weights and thresholds. The form of this representation constitutes the very design of the neural network, and holds the key to performance.

The subject of knowledge representation inside a neural network is very complicated, and becomes even more compounded when we have multiple interacting sources of information activating the network. Our present understanding of this subject is the weakest link in our knowledge of artificial neural networks. There are four common-sense rules for knowledge representation:

RULE 1. Similar inputs from similar classes should usually produce similar representations inside the network, and should therefore be classified as belonging to the same category. There are various mathematical methods for defining "similarity".

RULE 2. Items to be categorized as separate classes should be given widely different representations in the network. Rule 2 is the opposite of Rule 1.

RULE 3. If a particular feature is important, then there should be a large number of neurons involved in the representation of that item in the network. Actual needs depend on specific designs.

RULE 4. Prior information and invariances should be built into the design of a neural network, thereby simplifying the network design by not having to learn them. Rule 4 is particularly important because proper adherence to it results in a neural network with a *specialized (restricted) structure.* This is highly desirable for several reasons:

1. Biological networks are very specialized.

2. A neural network with specialized structure requires a smaller data set for training, learns faster, and often generalizes better.

3. Network throughput is accelerated.

4. The cost of building the network is reduced.[61]

Cost itself is not a philosophical issue, but its effects on neural networks may be: an issue for later discussion is the acceptability of a small artificial neural network representation of a large natural neural network structure.

Prior information may be built into the design of a neural network by using a

61. Haykin. *Ibid.* 24-26.

combination of two techniques: (1) *restricting the network architecture* through the use of local connections and (2) *constraining the choice of synaptic weights* by the use of weight sharing. The actual use of these two techniques in practice is strongly influenced by the application of interest. In general, the development of well-defined procedures for the use of prior information is an open problem.[62] Some built-in stability is desirable in order to maintain the usability of a trained network: further training on a different set of data may render a network less efficient at the task for which it was initially trained. Haykin refers to these stability issues under the rubric of "invariance to transformation."

There are three techniques for rendering classifier-type neural networks invariant to transformations: *Invariance by Structure*, *Invariance by Training*, *Invariant Feature Space*.[63] The use of an invariant-feature space may offer the most suitable technique for neural classifiers.

There is no well-developed theory for optimizing the architecture of a neural network which is required to interact with an environment of interest, or for evaluating the way in which changes in the neural network architecture affect the representation of knowledge inside the network. So we begin with a working notion of the aim of artificial intelligence:

> the development of paradigms or algorithms that require machines to perform tasks that apparently require *cognition* when performed by

62. Haykin. *Ibid.* 27.
63. Haykin. *Ibid.* 28.

> humans . . . we have purposely used the term "cognition" rather than "intelligence," so as to broaden the task tackled by AI to include perception and language as well as problem solving, conscious as well as unconscious processes.[64]

The system must have three specific capabilities:

(1) store knowledge;

(2) apply the knowledge stored to solve problems; and

(3) acquire new knowledge through experience. Obviously the brain performs these tasks, using a neural network.

Neural networks differ from classical AI in three specific categories:

(1) *Level of Explanation*. Classical AI models cognition as the *sequential processing* of symbolic representations.

Assumptions made in neural networks as to what explains cognitive processes are entirely different from those in classical AI. Neural networks emphasize the development of *parallel distributed processing* (PDP) *models*, which assume that information processing takes place through the interaction of a large number of neurons, each of which sends excitatory and inhibitory signals to other neurons in the network. Neural networks emphasize *neurobiological* explanation of cognitive phenomena.

(2) *Processing Style*. In classical AI, processing is sequential, as in typical computer programming. Operations are performed in a step-by-step manner. Parallel

64. Haykin. *Ibid.* 32-33.

processing is a distinctive feature of neural networks: parallelism is not only conceptually essential to the processing of information in neural networks, it is the source of their flexibility. Parallelism may be massive, giving neural networks a remarkable robustness. The automatic processing of contextual information is integral to neural networks. Knowledge is represented by the very structure and activation state of the neural network and not by declarative expressions. "The content necessary for a given problem is then nothing less than the whole neural network. Every neuron is potentially affected by the global activity of all other neurons in the network, with the result that context is dealt with automatically."[65] Although sequential operations may be simulated in parallel, and parallel operations may be simulated sequentially (as are many neural network simulations), the conceptual difference is both striking and important: it may well be that some cognitive activities are founded on emergent properties of the network that rely on simultaneity or close temporal connection of electrical or chemical activities. The levels of brain function may be compared to the levels of programming code, in which the source code of one level is data for the next level of processing, with the sodium ions taking the place of individual signals which activate some programmed function in the neuron which then combines signals with other neurons to send a signal to a local processing structure, which joins its signal with other local structures to feed the next level of processing,

---

65. Haykin. *Ibid.* 35.

and so forth.

(3) *Representational Structure.* Symbolic representations possess a *quasi-linguistic structure.*

For classical AI:

(1) Mental representations characteristically exhibit a combinatorial constituent structure and a combinatorial semantics.

(2) Mental processes are characteristically sensitive to the combinatorial structure of the representations on which they operate.[66]

In a neural network, representations are distributed, but it does not follow that whatever is distributed must have constituents, and being distributed is very different from having semantic or syntactic constituent structure. In a sense, the neural network goes beyond mere representation and itself becomes the information. The neural network responds systemically to the total environment. The network's ability to self-organize demonstrates the capability of a network to respond meaningfully to the environment without explicitly creating or manipulating symbols. Its representations are *subsymbolic*. It may be capable of computations beyond those of a Turing machine.[67]

In summary, we may describe symbolic AI as the formal manipulation of a language of algorithms and data representation in a *top-down* fashion. And we may

66. Haykin. *Ibid.* 35.
67. Seigelman, Hava. n.d. *Science*, n.d. 1994. per discussion with R. Miikkulainen.

describe neural networks as parallel distributed processors with a natural learning capability, and which usually operate in a *bottom-up* fashion. For the implementation of cognitive tasks, it may be more potentially useful to build *structured connectionist models* that incorporate both symbolic AI and neural networks. One may, in this way, combine the desirable features of adaptivity, robustness, and uniformity offered by

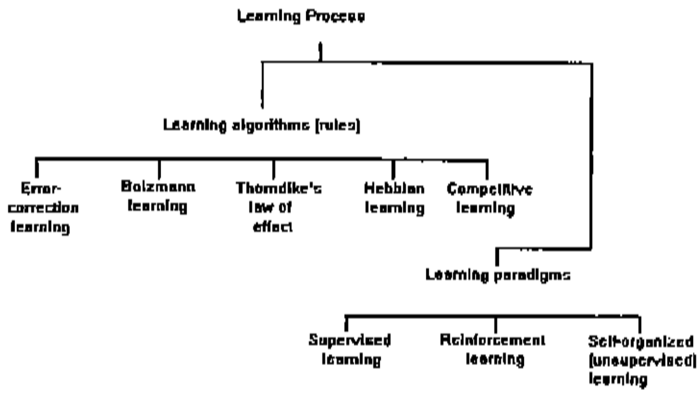

Figure 12. Haykin's Taxonomy of Learning Algorithms

neural networks with the representation, inference, and universality that are inherent features of symbolic AI. Such hybrid systems may be commercially useful, but a mix of approaches is sure to raise additional questions about the propriety of a particular simulation. The argument against getting semantics out of syntax resurfaces. Philosophy has traditionally considered the possibility of machine intelligence from a strictly symbolic viewpoint.[68] It is important to realize that a neural network operates in a radically different way than symbolic programs, and not all arguments regarding

68. See Appendix 3. The Chinese Room.

the logical limits of computation are valid in a neural network context.

Learning Processes[69]

Learning algorithms are sometimes based on biology, but are also sometimes rooted in mathematics. Haykin stipulates:

> *Learning is a process by which the free parameters of a neural network are adapted through a continuing process of stimulation by the environment in which the network is embedded. The type of learning is determined by the manner in which the parameter changes take place.*

The definition of the learning process implies the following sequence of events:

1. The neural network is *stimulated* by an environment.

2. The neural network *undergoes changes* as a result of this stimulation.

3. The neural network *responds in a new way* to the environment, because of the changes that have occurred in its internal structure . . .

A prescribed set of well-defined rules for the solution of a learning problem is called a *learning algorithm.* Neural networks are designed to take advantage of a diverse variety of learning algorithms, each of which offers advantages of its own. A *learning paradigm* refers to a *model* of the environment in which the neural network operates. These algorithms are based on mathematical models which may or may not represent similar activities in the nervous system.

69. See Appendix 5. Learning.

Neurobiological Considerations: Hebbian Learning

Hebb's postulate of learning has been a subject of intense experimental interest among neurophysiologists and neuropsychologists for many years.[70] Hebb's theory may be summarized as follows:

When an axon of cell A repeatedly and persistently takes part in firing cell B, a growth process or metabolic change occurs in one or both cells which results in an increase in A' s efficiency in exciting B.[71]

This model has been demonstrated to work in natural neural networks, and so is an important feature when implemented in artificial neural networks.

Competitive Learning

In *competitive learning*, the output neurons of a neural network compete among themselves for being the sole node fired. In a Hebbian neural network, several output neurons may be active simultaneously, but in a competitive learning neural network only a single output neuron is active at any one time. This feature makes competitive learning well suited to discover those statistically salient features that may be used to classify a set of input patterns.

The idea of competitive learning may be traced back to the self-organization of orientation sensitive nerve cells in the striate cortex. There is substantial evidence for

70. Hebb, Donald O. *The Organization of Behavior*. New York: Wiley, Introduction and Chapter 4, "The first stage of perception: growth and assembly," 1949. xi-xix, 60-78.
71. See Appendix 5. Learning.

competitive learning playing an important role in the formation of topographic maps in the brain.

Boltzmann Learning

The *Boltzmann learning rule* is a stochastic learning algorithm derived from information-theoretic and thermodynamic considerations. Even in this non-organically based example, the project of artificial neural networks does support our analysis at least in part of how the natural neural network functions: that any artificial neural network succeeds at all is at least a partial justification for continued investigation of their potential as working mathematical models of knowledge representation, whether brainlike or not.

Neurodynamical Models and Chaos

The brain is a non-linear dynamical system that lends itself to the study of chaos at several levels of function, and these chaotic dynamics also play an important role in the study of artificial neural networks, which can be used to model a chaotic time series.[72] In brain dynamics, it has been proposed that 1) chaos may provide the driving activity that is essential for Hebbian learning of novel inputs, 2) the long-term unpredictability of chaos may permit the brain to create new possible responses, suggesting a role for chaos in rapid adaptation to changing environmental conditions, and 3) sensitive dependence of chaos on initial conditions may provide an efficient

72. Haykin, *Ibid.* 550-551.

mechanism for dissipating perturbation. Neurodynamical systems of interest possess four general characteristics: 1) a large number of degrees of freedom, 2) nonlinearity, 3) dissipation, and 4) noise. Appropriately designed artificial neural networks and natural neural networks have these characteristics in common. In fact, the presence or absence of all these characteristics may be an important factor in determining the utility of a given artificial neural network as a model of natural neural network function.[73] [74]

Supervised Learning

Supervised learning requires an external teacher. The environment is unknown to the neural network, but it is provided with a set of input-output examples. when presented with a sample input from the environment, the teacher provides the neural network with an example of the desired response. An error signal represents the difference between the actual and the desired responses.[75] Such an artificial neural network represents essentially a back-propagation network with examples. The level of function is evidently high, and I currently know of no low-level functions it may represent. At a behavioral level, this design seems appropriate.[76]

Reinforcement Learning

If an action taken by a learning system is followed by a satisfactory state of

73. Haykin. *Ibid.* 551-552.
74. Gleick, James. *Chaos: Making a New Science.* New York: Penguin Books. 1988.
75. Haykin. *Ibid.* 57-59.
76. Dreyfus. *Ibid.* ¶ 60.

affairs, then the tendency of the system to produce that particular action is strengthened or reinforced.[77] Otherwise, the tendency of the system to produce that action is weakened.[78] "Although it cannot be claimed that this principle provides a complete model of biological behavior, its simplicity and common-sense approach have made it an influential learning rule."[79] This is behavioral emulation at a very high level.[80]

Learning Theory Summary

In general, learning theories used in neural networks may or not be supported by neurobiological evidence. In some cases, neural networks try very hard to mimic the structure and activity of the synaptic structures they represent in some detail. In other cases, the networks are designed to emulate higher level behaviors exhibited by biological cognitive systems, and in yet other cases, neural networks are designed to demonstrate the power of certain theories of learning which are purely conceptual, and which may or may not relate to actual brain functions. Issues discussed within the neural network community in this area will include whether the network design accurately models synaptic structures, but more often will focus on the relative benefits of various mathematical methods or their appropriateness to the particular model being examined. Neural networks have had some success in human-like tasks

---

77. Hilgard, Ernest R. *Theories of Learning.* New York: Appleton-Century-Crofts, Inc. 1956. 19-21. (Thorndike's *law of effect.*)
78. Haykin. *Ibid.* 59.
79. Haykin. *Ibid.*
80. Dreyfus. *Ibid.* ¶ 57.

such as training on both visual and sonic representations of words in one language, overlaid by training in a second language to produce a bilingual lexicon,[81] and in very non-human tasks such as differentiating between biological and mechanical objects discovered by SONAR signals. Self-organized mapping into a network applies to the brain -- especially evident in visual mapping -- and to artificial neural networks as evidenced in the DISCERN program.[82] In the example in Figure 12 (Miikkulainen's Figure 6), semantically related groups have been identified: the group including both *bat* and *chicken* seems to represent non-human animals; the group including *rock* and *window* are inanimate objects; note that *bat* may be either. This is an effective example of the power of networks to meaningfully self-organize.[83]

81. Soler, O. & R. van Hoe. "BAR: A Connectionist Model of Bilingual Access Representations", in *ICANN '94, Proceedings of the International Conference on Artificial Neural Networks*, Vol 1, Part 2. 263-267.
82. Miikkulainen, Risto. *Subsymbolic Natural Language Processing: An Integrated Model of Scripts, Lexicon, and Memory*. Cambridge: MIT Press, 1993.
83. Dreyfus. Ibid. ¶ 60.

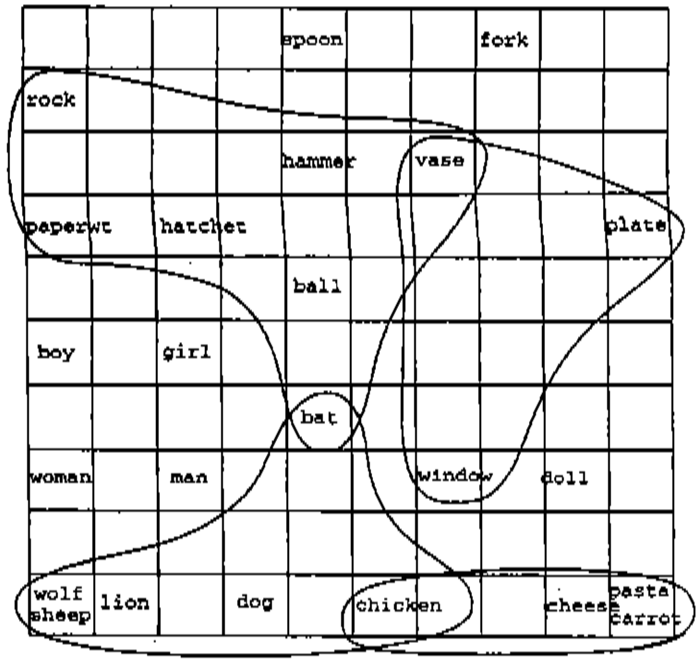

*Figure 6: **2-D Kohonen-map of the representations**. Labels indicate the maximally responding unit in the 10 x 10 feature map network for each representation vector. The map was formed in 15,000 epochs, where the neighborhood radius was decreased from 4 to 1 and the learning rate from 0.5 to 0.05 during the first 1,000 epochs, and to 0 during the remaining epochs.*

Figure 13. A typical self-organized network representation.[84]

84. Miikkulainen, Risto and Michael G. Dyer. "Natural Language Processing With Modular PDP Networks and Distributed Lexicon." *Cognitive Science*, 15 (3): 343-399, 1991. 11. Used by permission of author.

III. Philosophical Issues

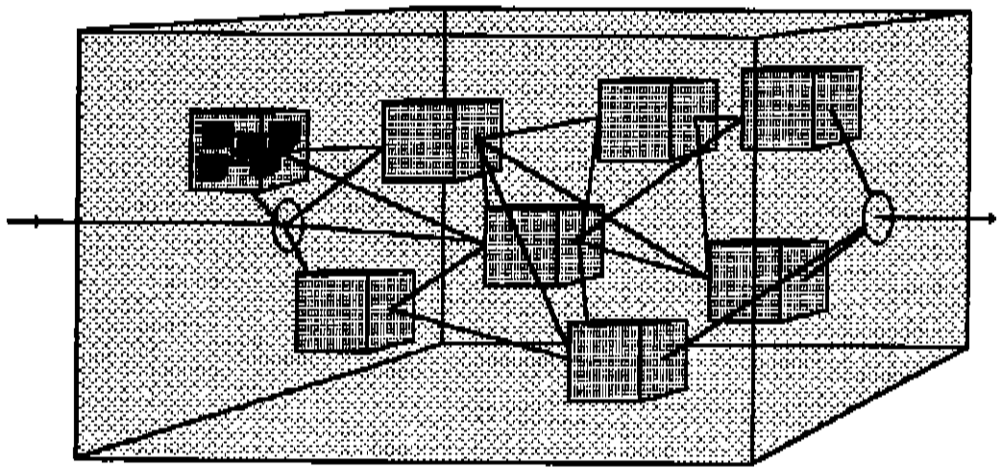

A Hierarchy of Black Boxes

Figure 14. Black Boxes within Black Boxes

The brain may be seen as a hierarchy of black boxes having, according to most sources, eight levels (Table 2). The analysis of any function or brain activity focuses on the comprehension of what exactly occurs at the level of interest. Whether the system can be usefully analyzed as a whole or must be decomposed into its component functions is an open question, as is whether any level may be understood in its own right without reference to the activity of its substrate. I will focus first on a comparison of natural and artificial neural networks, then on the current state of understanding of each. Then I will discuss the philosophical issues of interest. At the

Table 02. A sampling of possible neural network structures, three hierarchical and one relating reflex arcs at consecutive CNS levels

Haykin	Shepherd	Smolensky	Rahmann
Central Nervous System (CNS)	Behavior	Conceptual	CNS
Interregional Circuits	Systems and paths	Sub-conceptual	Telencephalon (cortex)
Local Circuits	Centers		Diencephalon (thalamus)
Neurons	Neurons	Neural	Mesencephalon
Dendritic Trees			Metencephalon (cerebellum, pons)
Neural Microcircuits	Microcircuits		Spinal cord
Synapse	Synapse		Periphery
Molecule	Membranes, molecules, ions		

neuronal level there is only a slight, but still remarkable, resemblance between the elements of an artificial neural network and the natural neural network. The basic principles of varying weights adjusted by training of synapses, multiple connectivity of nodes/neurons, plasticity of connection and activation level, are preserved. Despite massive differences in scale and incomparability of media, both networks perform certain activities in an apparently similar manner. The inputs to each system are quite different, even when an artificial network attempts direct acquisition of similar data. The television camera is nothing like the eye. A microphone is both more sensitive and less receptive than the ear. A robot arm can emulate touch through the use of pressure sensors, but what does it make of texture? There are automated odor sensors. Despite the fact that the operation of these electromechanical gizmos is nothing like the sensory apparatus of the human body, is it not plausible that an artificial neural network can make a representation of the data from these devices that is, in principle, not only as useable as that of the brain, but also is similar in internal representation? Can we not program such a system to have a sense of its own well-being? Given that the electromechanical systems available for such emulations are extremely limited, what are the issues of concern to cognitive science?

Hebb's principles seem to be born out in studies of human learning, and seem to operate as expected in an artificial network. Self-organization seems to work similarly in both media, naturally categorizing experiences in a meaningful way. Some learning algorithms seem to be shared, while others, such as Boltzman learning, seem to be exclusive to artificial networks. It is conceivable that both types of network operate solely on the principle of categorization. Some purposeful modeling is done in some artificial networks, and some structure is imposed by biology on natural networks. Otherwise, they seem to operate on similar principles, performing similar functions in the same way. The differences in scale imply no necessary difference in kind of operation. There is nothing mysterious in the tendency of inanimate matter to self-organize,[85] [86] and thus no contradiction when animate systems do so. Natural network memory relies on changes in synaptic potential[87] which may be modeled in artificial networks. The mapping of visual images in the brain[88] is very similar to a lexical distribution in a self-organizing artificial neural network.[89] Artificial neural network system outputs are limited to sound or print for the most part: fully mobile, self-directing operational neural-network-driven robots are not available. When developed, neural networks would give such entities far more flexibility than a symbolically-coded machine, without necessarily endowing it with humanity. We are very far from inducing a machine to share in the complexity of a natural neural network.

Symbolic and Connectionist Paradigms

Arguments between connectionists and symbolists as to which is the more appropriate and effective paradigm[90] seem to me to fall into the same consideration as arguments over whether one level of code is a program or simply grist for the program which is running it. At one level, or from one perspective, an issue may seem more compliant with one interpretation or another. Although it is certainly true that the

85. Umbanhowar. *Ibid.* 793-796.
86. Mukerjee, Madhusree. *Scientific American*, Vol 275, Nr. 5, November 1996. 28-36.
87. Rahmann. *Ibid.* 206-9.
88. Cook. *Ibid.* 32-39.
89. Miikkulainen and Dyer. *Ibid.* 11.
90. Dinsmore, John. *The Symbolic and Connectionist Paradigms: Closing the Gap* Hillsdale, New Jersey: Lawrence Erlbaum Associates, Publishers. 1992.

brain is capable of symbol creation, identification, and manipulation, this happens at a much higher level of mental function than the neuronal. What we have been able to show with artificial neural networks is that certain brain functions, such as categorizing and memory, may in fact, be emulated using a connectionist approach. The argument that these programs are themselves written in symbolic code is moot: if we follow the trail far enough, we will run in to the bottom end of quantum physics looking for the original program running the original data which forms the substrate for all else. In response to the argument that artificial neural networks perform tasks that no human normally does (i.e., SONAR identification of animate vs. inanimate objects) may be responded to with the fact that experiments with artificial neural networks may lead us to new routes of exploration in natural neural networks. As artificial neural networks become more complex and structurally representative, we should be able to investigate the operations of assemblies at the higher levels of mental operation. We are learning that the principles of representation by network are functional in both artificial and natural neural networks. Arguments that artificial neural networks don't succeed excellently with tasks like grammar[91] overlook the level of success they do have. The principle has been proven to work; it is a matter of refinement to get them to work more perfectly, to bring them to the same level of imperfection as human grammatical ability.

---

91. Pinker, Stephen. *The Language Instinct: How the Mind Creates Language*. New York: Harper Perennial, 1995. 132. I think Pinker is overly critical with the example of the irregular verb *smeej*. The artificial neural network-derived *leefloag* is only one letter short of the correct irregular past-tense.

There is a mix of non-overlapping and overlapping functions, capabilities, and structures in the artificial and natural networks. Similar structures seem to apply in representation. Palm *et alia* consider whether cell assemblies and neuron assemblies work alike in associative memory.[92] According to Rahmann and Rahmann:

> The neurons perform the specific, specialized activity of the nerve tissue, i.e., reception, processing, and transmission of information from cell to cell and, most importantly, the storage of information, thereby serving as the repository of memory content . . . a neuron, as the elemental unit of the nervous system, constitutes both an ontogenetic (developmental) and a physical (i.e., trophic and functional) unit that is responsible for perception, processing, transmission, and, above all, storage of information in an organism. [93]

Connectionists and symbolists argue over the workings of language: Searle's argument regarding syntax and semantics applies only to computational models, not neural network models. Searle comments that, "The reason that no computer program can ever be a mind is simply that a computer program is only syntactical, and minds are more than syntactical. Minds are semantical, in the sense that they have more than a formal structure, they have a content."[94] This assumes that the syntactic and symbolic elements are at the same level. They are at the same level in the

92. Palm, Günther and Ad Aertsen. *Brain Theory*. New York: Springer-Verlag. 1986.

93. Rahmann. *Ibid*. 2.

94. Searle, John. *Minds, Brains and Science*. Cambridge: Harvard University Press. 1984. 31. See Appendix 3: The Chinese Room. I have a vision of Searle and a large group of people in a connectionist Chinese room: instead of symbols coming in to a person/processor who uses a look-up to see what symbol to put out, a hand comes through the wall and taps one of those present on the shoulder. That person taps another on the shoulder, and so it goes until the output side of the room is reached. By this time, one or more of those near that wall will have been tapped on the shoulder, and they reach through the wall and tap others in other rooms on the shoulder. The process is not understood by those inside the room in this scenario any more than it is in the original Chinese room, but the point to be made is that this is not the appropriate level for conscious understanding. Searle's Chinese gym is no better an example.

computational model, but they are at different levels in a neural network. In a neural network, the structure *is* the content. Fischbach implies the failure of the symbolist paradigm to account for facial recognition, noting the implausibility of storing a symbol for one face per cell.[95] Somewhere in the hierarchy of layers, recognition happens. Somewhere in the hierarchy of abstraction, the network recognizes and manipulates symbols. Somewhere in the hierarchy of neuron connections, a person reacts to the outside world. Somewhere in the hierarchy of cell assemblies the mind is aware of at least some of its own content. Regardless of comparison to computer architectures, the brain obviously performs these functions, whether they be analyzed symbolically or as network. Point of view, purpose, context, level of operation -- all are factors in which approach is most appropriate.

Some network designs perform tasks not associated with brain neurons: the brain has few feedback loops from a neuron to itself; neurons do not switch between being excitatory and inhibitory. Some philosophical interest outside of the brain itself, such as possible modes of knowledge representation, are valid fields of application for artificial neural networks: that they work at all is of some significance for epistemology: somewhere between Socrates' "wax blocks" and Aristotle's "categories", neural networks may hold the keys to the truths hinted at by ancient speculation.

95. Fischbach. *Ibid.* 57.

(input(input(input(input(input(input process output)output)output)output)output)output)

(molecules)

(synapses)

(neurons)

(networks)

(layers)

(behavior)

Figure 15. A representation of the behavioral level making a call on the sublevels down through the hierarchy to the molecular level, the product of each level's computation being the result of its call to its substrate.

Implications of the Hierarchy

The hierarchical structure of neural networks presents interesting problems. We ask what it is at a substrate level that creates properties in the current level. Here again is a contrast between traditional AI and connectionist models: David Marr has proposed a three-level analysis of computational systems. According to Marr:[96] [97]

> 1. the computational level of *abstract problem analysis*, wherein the task (e.g., determining structure from motion) is decomposed into its main constituents;
>
> 2. the level of the algorithm, which *specified a formal procedure* by which, for a given input, the correct output would be given, and the task could thereby be performed;

96. Churchland, Patricia S. and Terrence J. Sejnowski. *The Computational Brain*. Cambridge: The MIT Press. 1992. 18-27.

97. See Appendix 6. Computational Autonomy.

> 3. the level of *physical implementation* of the computation.
> Higher levels are largely independent of the levels below it . . . hence, computational problems of the highest level could be analyzed independently of understanding the algorithm that executes the computation.

Can we apply Marr's three-level critique of algorithm implementation to the eight functional levels of the natural neural network? Marr's division treats computation as a single kind of level of analysis. Churchland and Sejnowski measure this system against the levels of organization of the nervous system and find it wanting. They list the nervous system levels as: molecules, synapses, neurons, networks, layers, maps, and circuits. Rahmann's list is connective rather than strictly hierarchical: Telencephalon (cortex), Diencephalon (thalamus), Mesencephalon, Metencephalon (cerebellum, pons), Spinal cord, Periphery. Rahmann's focus is memory, and he shows different circuit types for each level, some explicitly showing feedback loops.[98] The ultimate complexity of these structures in the multiplicity of functional levels is positively daunting.[99] Smolensky's three levels are neural, sub-conceptual and conceptual.[100] Computationally, we may see each of the levels as supplying data to feed the processor at the next higher level in the hierarchy. In some cases, this inter-level leap may seems like that between nuclear physics and chemical action -- or may simply seem like the leap between separate sentence fragments and a complete sentence. If we analyze this structure as a computational structure, should it

98. Rahmann. *Ibid.* Fig. 5.6. p. 106.
99. See Appendix 1: Memory.
100. Smolensky, Paul. "On the proper treatment of Connectionism." *Behavioral and Brain Sciences* (1988) 11. 1-74.

not be represented simply as IPO: input, process and output? In a symbolic program, the output of a subroutine or function is the result of a combination of its own processing and that of the substrate of functions and subroutines on which it calls. In theory, this should also be true of a connectionist program.

Does this theory take into account the character of a systemic response? In a neural network, levels 2 and 3 are combined: the electrobiophysical structure of the network is the embodiment of the algorithm -- formal or not, and it would appear that abstract problem analysis is a very high level activity.

If we are to look at three levels, should they not simply model *knowledge representation*: where is the "A" in my brain when I *see* "A & B $\rightarrow$ C"? If I train the net to read characters, the symbols "O" looks like 0.01 0.85 0.76 . . . etc. What I do with information at the next level is likely to be similar in both representations, but where is a picture of the next level?

The usual concept is that a picture of "A & B $\rightarrow$ C" is decomposed into separate parts, then reintegrated. This would represent Marr's Level 1, which certainly applies to formal problem solving. This decomposition usually doesn't happen. Usually "o" is seen as in "m*o*del," "h*o*t" -- as one element of a total picture -- we may examine letters in order from left to right (in English), but first we see a whole word or phrase -- model is not read "m," "o," "d," "e," "l." It is read "model."[101] We only

101. Dennison, Paul E. & Gail E. Hargrove. *Personalized Whole Brain Integration.* Glendale, CA: Edu-Kinesthetics, Inc. 1985. 9.

examine it letter by letter if there is sufficient ambiguity in the context so that it needs to be differentiated from "motel," "modal," "Gödel," etc. To insist on a constructive action in every act of reading is like insisting on the analysis of a picture pixel by pixel beginning in the upper left corner. That would be like looking at the Mona Lisa and seeing it not as one picture, but as a collection of related shapes which happen to begin at the upper left corner of the picture frame. We see first a picture of a page, then of a word, then, if necessary, of individual letters. (Smolensky would include a syllabic intermediate form.)[102] Its parts may be individually identified, but we had to learn at one time to concatenate those previously learned parts and to construe their adjacency in a certain way -- according to certain model. The neural model which *identifies* "A" interacts with other models (concatenation, for example). One interpretation of *intelligence* is: the ability to effectively identify and manipulate appropriate models. I see the phrase "A & B $\rightarrow$ C". When I first take note of something, I immediately try to contextualize it. (Much of humor relies on ambiguous contextualization.) As I gather information about the thing-noticed, I focus on the context. I may first see "A & B $\rightarrow$ C" as 'some words on a page' -- this context is immediately apparent from the appearance and location of the statement. Then I may take note of its structure and see it as a 'phrase or sentence.' When I note the character content, I see that it is a 'logical statement'. Once I have identified the

102. Smolensky.

contents in such a way as to accurately contextualize it, I may begin to apply "rules" associated with the identified context as they apply to the particular content.[103]

This decomposition is sometimes necessary, but in the course of daily events, the context usually gives me all I need to know to process the statement with a cursory glance at a couple of its details. I may notice first the structure of the sentence -- a conjunctive implication, then the order of the variables used. I can write an algorithm for a symbolic program that will analyze such a statement one character at a time and report on whether it is well-formed, or whether it is justified by implication from other logical statements. Whether a symbolic program can be written that will tell me anything about the application of the statement depends on the richness of the context in which the statement applies. In a simple model, a simple symbolic program may do, but in the rich context of the multiply-modeled real world, a comprehensive application is highly improbable. The neural network picks up on all cues within its input domain. Whether they are the cues we consciously think are important does not play a role. So in similar circumstances, a neural network is responding to not just narrowly defined inputs which conform to certain symbols within a prescribed range, but also includes peripheral information, perhaps meta-information which may or may not apply effectively in achieving a desired result. Even when a conscious being makes a determined attempt to focus on what are

103. See Appendix 3. The Chinese Room.

considered the pertinent factors in solving a problem, his neural network may bring unsolicited (unconscious) factors to bear. The formal account of the brain's function at any given level is incomplete. That being the case, it seems impossible to insist that substrate definitions and activities are of no importance. The neural network is whole. Its hierarchical structure is integral to its function. It is a vertical as well as a horizontal program. It responds as a single organ.

When neural networks respond, we don't always know the rationale behind their behavior. Even we conscious creatures may not be aware of all the inputs to a decision: *intuition* is an issue for neural network study: how does a neural network process non-verbal models? What happens intuitively or reflexively? What is represented and how, and how do we govern behavior which is not the result of careful consideration? To what extent is consciousness of the mind's content necessary for responsibility? If chaos theory *does* represent a model active in the operation of the brain, at what level is it significant? Is there a level which is demonstrably non-deterministic? Is there a valid sense in which unpredictability implies a caesura in causation? Are there truly emergent properties in the hierarchy of structures, or will an adequately comprehensive scientific explanation imply all possible behavior, both animate and inanimate. Peter Clark[104] cites the demise of arguments such as William Thompson, First Baron Kelvin's partition of laws for living

104. Knowles, Dudley, ed. *Explanation and its Limits*. Cambridge: Cambridge University Press. 1990. 166-168

and dead matter, and supposes arguments regarding 'emergent properties' may go the same historic way to obsolescence. Until new explanations are developed to account for currently unexpected phenomena, for example 'oscillons,'[105] it may be beneficial to speak of unexplained properties as 'emergent.' This restricted use of the term -- to refer to the observed, but unexplained outcomes of systems -- seems to me to be both concise and appropriate. I expect that we will develop more comprehensive theories that render the notion of 'emergent' as moot, but in the interim, it is a handy term.

> The expression "Computational neuroscience" reflects the possibility of generating theories of brain function in terms of the information-processing properties of structures that make up nervous systems. It implies that we ought to be able to exploit the conceptual and technical resources of computational research to help find explanations of how neural structures achieve their effects, what functions are executed by neural structures, and the nature of representation by states of the nervous system.[106]

I worry about the tendency of cognitive scientists to use phrases referring to "states" of a neural network. Our natural neural network is in flux from the time it first begins to form to the time its operation comes to a halt. A living brain is a *dynamic* brain, *not static*. The only real-life "state" of a brain is death. It is sometimes helpful to use the *fiction* of a "state of the nervous system" in order to have a stable platform from which to launch a discussion, but the temptation is great to take this fiction as a possible *fact* of a natural neural network, and to overlook the dynamic

105. Umbanhowar. *Ibid.*

106. Churchland, Patricia, Christoph Koch and Terrence J. Sejnowski. "What Is Computational Neuroscience?" in Schwartz, Eric. *Computational Neuroscience*. Cambridge: the MIT Press. 1990.

implications for a truly effective artificial neural network simulation.

> [Computational neuroscience] also connotes the potential for theoretical progress in cooperative projects undertaken by neurobiologists and computer scientists. This collaborative possibility is crucial, for it appears that neither a purely bottom-up strategy nor a purely top-down strategy for explaining how the brain works is likely to be successful. With only marginal caricature, one can take the purely bottom-up strategy as recommending that higher-level functions can be neither addressed nor understood until all the fine-grained properties of each neuron and each synapse are understood. But if, as it is evident, some properties are *network* effects or *system* effects, and in that sense are *emergent* properties, they will need to be addressed by techniques appropriate for higher levels and described by theoretical categories suitable to that level. Assuming there are system properties or network properties that are not accessible at the single-unit level, then knowing all the fine-grained detail would still not suffice to explain how the brain works.[107]

An example: suppose one argued that it is necessary to know the exact properties and function of each muscle cell in the arm before being able to determine anything about the function or action of the whole arm or of the whole muscle of which a particular cell is a part -- this would be considered quite foolish. If one considers the brain and its constituent cells to have the same type of aggregate abilities in performing its overall function that a muscle and its constituent cells has in performing its overall function,[108] one may see that there are multiple levels of description of brain function, each of which may bear some fruit in leading to an understanding of the whole. Organs are organic in their activity. Each of their constituents works with the total system to perform the functions of the whole organ,

107. Churchland, et al. *Ibid.* 46.
108. Shepherd. *Ibid.* 83.

itself a constituent of the whole body. (note: cutting off a leg does not immediately affect the structures of the brain, hence phantom pain and continued sensation after an amputation.)

> A purely top-down strategy is typified, again with minimal caricature, by its dismissal of the organization and structure of the nervous system as essentially irrelevant to determining the nature of cognition. Advocates of this strategy prefer instead to find computational models that honor only (or at least primarily) psychological and computational constraints. One major reason for eyeing skeptically the purely top-down strategy is that computational space is consummately vast, and on their own, psychological and engineering constraints do not begin to narrow the search space down to manageable proportions. Unless we go into the black box, we are unlikely to get very far in understanding the actual nature of fundamental cognitive capacities, such as learning, perceiving, orienting and moving in space-time, and planning.[109]

In spite of great advances and continuing refinement in our understanding of natural neural networks, there are a great many things we still don't know about the facts of their operation. What controls their formation, structuring and restructuring? How does one level become the source of unexpected properties in the next hierarchical level? Can we develop a transitional calculus which can predict all of the next level's properties? Is this knowledge essential to meaningful discussion? Is knowledge reducible to a few basic tactics of categorization (or of categorical organization)?

Eric Schwartz[110] asks a series of questions, offering alternative responses:

c63E1. What is the relevance of simple model neural systems?

109. Churchland, et al. *Ibid.* 5.

110. Schwartz, Eric L., ed. *Computational Neuroscience.* Cambridge: MIT Press. 1990. xi.

> *Position 1*: Mammalian nervous systems are far too complex to study the properties of neural networks. We must begin with the simplest model system available, either in vertebrate organisms or simple localized synaptic modules in vertebrates.
>
> *Position 2*: Two decades have gone by, and the early rhetoric of model-systems proponents has not been fulfilled. Brute-force analysis of simple systems will not lead anywhere. Invertebrate behavioral capabilities and neural properties are remote from those of vertebrates. Simple vertebrate systems are of interest only in themselves, not as general brain models. Synaptic-level studies have not provided much insight into computational function.

I have to agree that simple models are the place to start, but if we want to understand the complex operations of the brain, we must strive to reflect the brain's connective complexity. The many types of circuitry typical of the various levels of the nervous system's hierarchy are indicators that a simple model will at best emulate no more levels of function than the number of layers in the design.[111]

> 2. Is physical locality an essential part of neural manipulation?[112]
>
> *Position 1 (homage to the Turing machine)*: The anatomical structure of the brain has no more to do with its function than the shape of the cabinet of a VAX, or the location of its circuit boards. Brain function is determined by the logical and dynamic connection properties of its neurons. The actual physical structure, location, architecture, and geometry is irrelevant compared to its logical, connectionist aspects. One could take a brain and grossly deform the position of its neurons, keeping only the topology of connections intact, and there would be negligible difference in performance.
>
> *Position 2 (computational anatomy)*: Significant recent work in brain research has been related to the discovery and elucidation of detailed forms of somatotopic mapping, laminar specialization in cortex, and columnar architectures representing sensory submodalities. These forms of functional architecture may represent a major mode of brain function: the formatting of sensory data in a manner that

111. Rahmann. *Ibid.* 106.
112. Schwartz. *Ibid.* xii.

> simplifies its further processing. One of the major differences in computational style of brain versus VAX may well be the indifference of the VAX to its geometry and the exquisite attention paid by the brain to its geometry.

I am not yet ready to give up on the ability of the artificial neural network to adequately represent the complete structure and activity of the brain, but the physical proximity and spatial structure of connecting neurons must surely be a vital consideration for realistic emulation. So too must be considerations of the medium and motivating force behind the operations of the network: the brain has neurotransmitters which act as a mass biochemical input which has systemic effects. Whatever governs the density of the neurotransmitter at each synapse must be of systemic significance, and this factor is not adequately represented in artificial neural networks. The artificial neural network system of weights and corrections is itself a step in this direction, and it may be malleable enough to form a basis for emulating the neurotransmitter background for neural function.

> 3. What is the appropriate balance of theory and experiment in neuroscience?
>
> *Position 1*: Theoreticians have constructed models which have little connection to experimental data, and which provide little opportunity for experimental test. Experimentalists will learn whatever is necessary to perform their work. So far, they have not had to bother with learning theory (mathematics, computer science, . . .). Theoreticians will have to do their homework, i.e. learn something about experimental disciplines which they are modeling.
>
> *Position 2*: Experimentalists have too little background to appreciate or understand the relevance of theoretical work. A science with no theoretical component is just a mass of phenomenological details, a form of "butterfly collecting." Experimentalists will have to do their homework, i.e. learn something about the theoretical tools

which are essential to their disciplines.

Obviously, everyone needs to do their homework. There are considerations pertinent to the ongoing search for meaningful results in cognitive science, not just from the traditionally interested fields of psychology, philosophy, neurobiology, linguistics and computer science, but also from education, physics, music[113] and religion. The greater the correlation of all these fields in a neural network context, the more likely a neural network modeler is to developing appropriate widely-applicable models. Focus on a very narrow field must be balanced occasionally by global considerations and searches for related models in other systems. For example, I'm sure the development of Teflon® was fascinating, but until someone related it to frying an egg, it was of limited significance to humanity at large.

> 4. What is the correct choice of spatial scale for modeling the computational abilities of brains?
>
> *Position 1 (synaptic-neuronal level)*: Since the brain is composed of neurons, whose properties are accessible to study and fairly well understood, the neural scale is the correct one at which to approach brain computation. Brains are networks of neurons, and the mathematical properties of such networks determine brain function.
>
> *Position 2 (column-map level)*: The mathematical tools available for studying large networks are woefully inadequate. The two-body problem in mechanics is easy; the three-body problem is very hard. The neural N-body problem presents difficulties vastly greater than those of classical statistical mechanics, a field which is largely computationally intractable. Neural firing densities need to be averaged into "densities," and the large-scale (i.e. columnar, map) properties of these densities made the basis for study. Just as fluid mechanics began with a simplified continuum hypothesis, neural modeling must find a simplified continuum level in order to "get off the

113. Scarborough, Don L., Ben O. Miller, and Jacqueline A. Jones. "Connectionist Models for Tonal Analysis." in Todd et al (q.v.) *Music and Connectionism*. 1989.

ground."

Since the brain is composed of neurons in neurotransmitter bath, it may be that the neural scale is only an approximation which can be refined. A systems model will be woefully inadequate if we haven't modeled the constituents correctly. We should be looking beyond the neural level in both directions. There are significant outcomes to be found both in looking more closely at the neurotransmitter input to the neural substrate and at the system effects of the network as a whole.

> 5. What is the correct choice of temporal scale for modeling neural computation?
>
> *Position 1*: The brain may be viewed as a dynamical system. Stable states of neural networks may be reached, since the time constant of neurons is in the range of 1 - 10 milliseconds, and differential equation systems have been shown which reach their stable states in only a few iterations.
>
> *Position 2*: Pre-attentive perception refers to the period about 200 milliseconds after an event. Most perceptual and many cognitive functions can occur in roughly this time period. But 200 milliseconds is also roughly the time of transit for a signal through the brain. It appears that the brain is a one-cycle machine, something like a lookup table! There is no time for any settling into stable states. Moreover, there is no experimental evidence that the brain ever "settles" into anything like an equilibrium condition.

Of course not! It would be like saying that a pulse can go through the brain, leaving no activity in its wake, and that no other pulse follows it immediately -- the patient must be dead! The necessity to freeze an artificial neural network at some point, or to limit its possibilities for further training[114] in order for it to retain its usability is an obvious departure from accurate emulation of brain function.

114. Haykin. *Ibid.* 27.

> 6. Is the contemporary Pavlovian position correct?[115]
>
> *Position 1*: Perception cannot be studied in isolation from concomitant motor or other goal-directed activity.
>
> *Position 2*: One can show, for example, a Julesz random-dot stereogram to someone who has never seen one before, and this person will see a stereo percept, without even knowing what he is looking at. Perception can be isolated, both functionally and computationally, from goal-directed activity.

For me, the question of perception always entails questions of consciousness. a neural net may be trained to 'see,' but is that the same as 'perceiving'? This paper does not have scope enough to consider consciousness in detail, and so I will forgo details discussions of consciousness and perception. It seems reasonable to say that without some form of feedback, we can't tell if a system, natural or artificial, has appropriately perceived something.

> 7. Are conventional computer metaphors valid for the brain?
>
> *Position 1*: It is crucial to distinguish algorithm from implementation in AI and neural modeling (David Marr). The algorithm need have nothing to do with the under-lying physical structure of the neuronal "hardware."
>
> *Position 2*: The implementation (brain structure) depends crucially on the nature of the computation (brain function). In the brain, the medium is the message.

I firmly believe that Marr's analysis is of considerable importance in a symbolic programming milieu. In a neural network context, however, it lacks comprehension. The systemic nature of the work of the network is not capturable in a purely symbolic program context. Marr's analysis applies to systems with a clear-cut syntax where the

115. Schwartz. *Ibid.* xiii.

syntactic level is in fact where the work of the system is done. The syntactic level of a neural network is at a much higher level of abstraction than the level where the work of the system is done. As we saw earlier, the comprehension of "A & B $\rightarrow$ C" is holistic in a neural network: the picture is decomposed into symbolic elements only when contextualization is inadequate. The picture is analyzed whole, not taken apart and then reassembled.[116] [117]

Real Neuron Architectures

Shepherd's considerations of "The Significance of Real Neuron Architectures for Neural Network Simulations" include this quote from John von Neumann:

> "One may have to face situations in which there are, say, hundreds of synapses on a single nerve cell, and the combinations of stimulations on these that are effective are characterized not only by their number but also by their coverage of certain special regions on that neuron (on its body or on its dendritic system) by the spatial relations of such regions to each other, and by even more complicated quantitative and geometrical relationships that might be relevant."

Shepherd concludes that

> "all complications of this type mean, in terms of the counting of basic active organs as we have practiced it so far, that a nerve cell is more than a single basic active organ, and that any significant effort at counting has to recognize this. If the nerve cell is activated by the stimulation of certain combinations of synapses on its body and not by others, then the significant count of basic active organs must presumably be a count of synapses rather than of nerve cells."

116. Hopfield, J. J. "Neural networks and physical systems with emergent collective computational abilities." *Proceedings of the National Academy of Sciences* 79:2554-2558, 1982. "A study of emergent collective effects and spontaneous computation must necessarily focus on the non-linearity of the input-output relationship. The essence of computation is nonlinear logical operations . . . Those neurons whose operation is dominantly linear merely provide a pathway of communication between nonlinear neurons."
117. Shepherd. *Ibid.* 83.

For Shepherd, this establishes the synapse as the basic computational unit of the nervous system.[118] The question of what drives the system remains outside most implementations of artificial neural networks. Kosslyn and Koenig focus on the relation of emotion and cognition: "We have seen that there is no reason to consider emotion and cognition as disparate events, and many reasons to view the two as inextricably bound. Emotion not only plays a role in prioritizing goals, but also is realized as a profile of which goals and processes are facilitated and which are inhibited."[119] Neural networks provide a means of processing sensory inputs and of retaining memory.[120] They may form in their connections the structure of the very thoughts we are aware of thinking. The relation between the network itself and its systemic controller is still an area needing extensive investigation. The question of motivation has always been a puzzle for psychologists; now it is an open question for all of cognitive science.[121]

Conclusions

Which black box can be safely ignored for philosophical purposes? At which level can one find philosophically significant implications? Much depends on what research, especially in natural neural networks, produces. I think each level may be profitably considered both in isolation and in its relation to the other levels. There is

118. See Appendix 4. Defining the Basic computational Unit.
119. Kosslyn, Stephen M. and Olivier Koenig. *Wet Mind: The New Cognitive Neuroscience*. New York: The Free Press. 1995. 471.
120. See Appendix 1. Memory.
121. Dreyfus. *Ibid.* ¶ 33.

always interest in trying to get to the bottom causal level, so substrates will continue to be worth discussing. The rules consistent within a level are evidence of what is and what happens. Maybe there are clues to ontological questions in the structure and activity of the network. The behavior of the total system may be an indicator of what is of value to the system. This valuation is of interest to Ethics.

Philosophy of mind is probably the most pertinent philosophical department to these issues. Thought, consciousness, perception, cognition, motivation, all are related to the neural network paradigm.

Philosophy of language certainly may benefit from experimental findings on the way people really create discourse.

Epistemology[122] will be enriched by the neural network paradigm for knowledge representation and perceptual processes.

The results achieved so far in artificial neural networks show that our interpretation of natural neural networks is at least one valid approach. Therefore, continued research in artificial neural networks is likely to produce meaningful results applicable to natural neural networks.

As one possible real-world example of how cognition occurs, neural networks cannot be overlooked when discussing any epistemological topic. Therefore, any serious discussion of what may be known by whom and how they may know it, must

---

122. O'Connor, D. J. and Brian Carr. *Introduction to the Theory of Knowledge*. Brighton, Sussex: The Harvester Press. 1982.

address the neural network interpretation.

Neural networks have proved themselves as models for knowledge representation. Therefore, any discussion of computational knowledge processing which reaches beyond simple symbolic representations of data type, field, record, etc., must acknowledge the efficacy of self-organizing systems, trainable perception devices, and black-box memories whose internal structure is itself the knowledge representation.

Experiments on artificial neural networks may lead to a better understanding of natural networks, by showing experimentally which structures are viable and which not. They may even indicate possible teaching techniques, psychological evaluations and psychiatric therapies.

Development of understanding neural network system dynamics may lead to a transitional calculus which may allow us to predict what now appear to be emergent properties.

We may discover that the neural network is only biology's hand-maiden -- that the essential us is in the biochemical reactions which drive natural neural networks.

"Spinoza . . . has said, 'we examine the universe through the lens of philosophy.'"[123] Neural networks may help us to raise the power of the lens and bring it to a finer focus.

---

123. Darden, Severn. "Metaphysics Lecture." *The Sound of My Own Voice, and Other Noises*. Mercury LP OCS 6202. 1961.

Appendix I. Memory

Rahmann describes the workings of memory:

> The functions of the neuroglia, in contrast to those of the neurons, cannot be defined so readily. Indeed, the glial cells insulate, protect, and support the nerve cells from external, mechanical influences, but the neuroglia also perform metabolic tasks of support and assistance in the sense of metabolic symbiosis with the nerve cells. Certain impulse-conducting properties of the neuroglia as well as tasks of support and assistance in neuron differentiation must be considered also.[124] . . . a neuron, as the elemental unit of the nervous system, constitutes both an ontogenetic (developmental) and a physical (i.e., trophic and functional) unit that is responsible for perception, processing, transmission, and, above all, storage of information in an organism. [125] Synaptic plasticity, in the sense of an experience-based change in neuronal function, can be established utilizing methods borrowed from the area of electrophysiology insofar as neurons are able to change their *bioelectrical response behavior* relative to stimulation. A presynaptic action potential effects changes in postsynaptic potential through the event of stimulation. Two categories of presynaptic mechanisms can account for this. In the first instance, it is the activity of the presynaptic terminal itself that causes a more or less short-term, stimulation-dependent change in the post-synaptic cell. In the second case, long-term changes in synaptic function are determined by the effect of modulator substances in the area of the synaptic contact zones.[126] . . . synaptic plasticity in neuronal networks has been documented in the form of substantially prolonged *adaptations in bioelectric responses* in the nervous system as well as in long-term *adaptive changes* in higher associative performance in the *process of learning and memory* that parallel those prolonged adaptations.[127] The ability of higher vertebrates to form abstractions and generalizations might well be the basis for all learning processes. The formation of conceptual complexes (i.e., abstract, nonverbal concepts) is one result of this ability. These are not characterized by symbols such as words . . . the ability to form a *concept of similarity* has been

124. Rahmann. *Ibid.* 1.
125. Rahmann. *Ibid.* 1-2.
126. Rahmann. *Ibid.* 2.
127. Rahmann. *Ibid.* 206.

corroborated[128] . . . In all of these learning processes, particular significance is placed on the *predisposition* of the individual (i.e., on the genetic foundation, on early perinatal experience) and on motivational conditions (positive or negative feeling relative to learning). [129]

According to Fuster:

"Long-term potentiation (LTP) in the mammalian brain is the outstanding electrophysiological phenomenon of persistent change in synaptic strength as a result of impulse transmission across synapses. Viewed by many as the biophysical basis of Hebb's postulate, that phenomenon has been extensively investigated in the hippocampus, an ancient cortical structure with simple and well-understood connectivity."[130] Nevertheless, as of 1995, Fuster still confesses that "the cellular mechanisms by which the repeated impulse conduction from cell A to cell B may lead to permanent change are not yet precisely known."[131]

Dreyfus takes a different approach:

[48] Neural networks provide a model of how the past can affect present perception and action without needing to store specific memories at all. It is precisely the advantage of simulated neural networks that past experience, rather than being stored as a memory, modifies the connection strengths between the simulated neurons. New input can then produce output based on past experience without the net having to, or even being able to, retrieve any specific memories. The point is not that neural networks provide an explanation of association. Rather they allow us to give up seeking an associationist explanation of the way past experience affects present perception and action.[132]

---

128. Rahmann. *Ibid.* 209.
129 Rahmann. *Ibid.* 241.
130. Fuster. *Ibid.* 25.
131. Fuster. *Ibid.* 25.
132. Dreyfus. *Ibid.* ¶48.

Appendix 2. Emergent Properties.[133]

In physical systems made from a large number of simple elements, interactions among large numbers of elementary components yield collective phenomena such as the stable magnetic orientations and domains in a magnetic system or the vortex patterns in fluid flow. Do analogous collective phenomena in a system of simple interacting neurons have useful "computational" correlates? For example, are the stability of memories, the construction of categories of generalization, or time-sequential memory also emergent properties and collective in origin? [Hopfield] shows that important computational properties spontaneously arise.

. . . Perceptrons were modeled chiefly with neural connections in a "forward"direction . . . made random net of neurons deal directly with a real physical world and did not ask the questions essential to finding the more abstract emergent computational properties . . . Perceptron modeling required synchronous neurons like a conventional digital computer. There is no evidence for such global synchrony and, given the delays of nerve signal propagation, there would be no way to use global synchrony effectively. Chiefly computational properties which can exist in spite of asynchrony have interesting implications in biology.

[Hopfield's] model . . . will use its strong nonlinearity to make choices, produce categories, and regenerate information and, with high probability, will generate the output . . . from such a confusing mixed stimulus.

Most neurons are capable of generating a train of action potentials--propagating pulses of electrochemical activity--when the average potential across their membrane is held well above its normal resting value. The mean rate at which action potentials are generated is a smooth function of the mean membrane potential . . .

A study of emergent collective effects and spontaneous computation must necessarily focus on the non-linearity of the input-output relationship. The essence of computation is nonlinear logical operations . . . Those neurons whose operation is dominantly linear merely provide a pathway of communication between nonlinear neurons.

. . . The neural architecture of typical cortical regions and also of simple ganglia of invertebrates suggests the importance of 100-10,000

133. Hopfield, J. J. "Neural networks and physical systems with emergent collective computational abilities." *Proceedings of the National Academy of Sciences* 79: 2554-2558. 1982.

> cells with intense mutual interconnections in elementary processing . . . In the model network each "neuron" has elementary properties, and the network has little structure. Nonetheless, collective computational properties spontaneously arose. Memories are retained as stable entities or Gestalts (sic) and can be correctly recalled from any reasonably sized subpart. Ambiguities are resolved on a statistical basis. Some capacity for generalization is present, and time ordering of memories can also be encoded. These properties follow from the nature of the flow in phase space produced by the processing algorithm, which does not appear to be strongly dependent on precise details of the modeling. This robustness suggests that similar effects will obtain even when more neurobiological details are added.
> Much of the architecture of regions of the brains of higher animals must be made from a proliferation of simple local circuits with well-defined functions. The bridge between simple circuits and the complex properties of higher nervous systems may be the spontaneous emergence of new computational capabilities from the collective behavior of large numbers of simple processing elements.
> Implementation of a similar model by using integrated circuit chips would lead to chips which are much less sensitive to element failure and soft-failure then are normal circuits. Such chips would be wasteful of gates but could be made many times larger than standard designs at a given yield. Their asynchronous parallel processing capability would provide rapid solutions to some special classes of computational problems.

Peter Angeles' *Dictionary of Philosophy* offers:

> **emergent, an.** Sometimes *gestalt property* of organized structures. The new qualitative synthesis produced by structures organized in certain patterns that cannot be predicted from the examination of the constituent parts of the whole.[134]

134. Angeles, Peter A. *Dictionary of Philosophy.* New York: Barnes and Noble Books. 1981. 73.

Appendix 3. The Chinese Room.

"The Chinese room shows what we knew all along: syntax by itself is not sufficient for semantics."[135]

Searle is not wrong in drawing this conclusion from what we know about syntactic models. He seems to persist in assuming, though, that all computation is based on symbolic processing, including neural network computation. He assumes that at the heart of a neural network there is "a set of rules, the program," which manipulates symbols, whether programmed as such, or as in the particular symbolic/neural hybrid proposed by Harnad, "took in stimuli, [and] converted these into symbols . . ." His argument against the System Reply specifies "In order to justify the System Reply one would have to show: How the system gets from the syntax to the semantics . . ." There is no syntax in a neural network; it's operation is semantic to begin with. The rules a neural network learns are not built up from elemental individual rules, like a formal grammar, but are molded out of a total picture in which individual elements are weighted by their apparent affect on the whole. He continues: "Behavior plus syntax is not constitutive of cognition. To repeat, where the ontology --- as opposed to the epistemology of the mind is concerned, behavior is irrelevant . . . If we define the nets in terms of their computational properties, they are subject to the usual objection. Computation is defined syntactically and syntax by itself is not sufficient for mental contents. If we define the nets in terms of physical features of their architecture then we have left the realm of computation and are now doing speculative neurobiology. Existing nets are nowhere near to having the causally relevant neurobiological properties." I agree that no artificial current network is anywhere near ample for a full-blown emulation of a complete natural neural network. His argument that syntax is inadequate for semantics misses the point of neural networks -- they are not symbolic programs at the level of operation, even though they must of necessity be programmed in a syntactic-model programming language. His focus is on the objection to a set of rules. somewhere, at the bottom of things, if nowhere else, is a set of rules which govern all that is. These rules may be quantum-mechanical or determinist-causal rules, but whatever they are, they are consistent and universally applicable. some application of these rules results in cognition, whether this be considered as conscious thought or as behavior or as input-process-output computation. I don't know that these rules conform to a notion of symbolic manipulation; they are rules which constitute being and eventuality -- they are, if anything here may apply, semantic. They apply as a single force, not as a composition of symbols. The rules rule!

People speak in complete semantic units before they learn to analyze those

135 Searle, John R. "The Failures of Computationalism." (Response to Stevan Harnad (q.v.), "Grounding Symbols . . .") *Think Quarterly*, Vol. 2, No. 1, 1993. Instituut voor Taal- en Kennistechnologie (http://tkiwww.kub.nl:2080/tki/Docs/Think)

units into the component parts of a syntactic model. given that a neural network does not need syntax in order to perform semantically, how can we usefully differentiate whether *any* system is a cognitive system? What are the criteria for cognition? The Chinese room does show that passing the Turing test is no guarantee: a symbol processing system could conceivably pass the test. Is we substitute a properly trained neural network for the symbolic program, then it depends on whether our definition of cognition will be satisfied by observable behavior or be satisfied only by evidence of consciousness. If the latter, then ask how we determine consciousness. Is any system which modifies its behavior in response to feedback from its environment conscious? Do we require a system to exhibit a particular observable behavior, that of a self-report of consciousness. Is it enough that it exhibits behavior we attribute to consciousness, or will we hold out for a verbal declaration of self-awareness? we know that a proper application of the fundamental rules yields consciousness. would we exhibit this interested behavior if we were not ourselves conscious? perhaps my claim to consciousness is merely a behavioral outcome of my neural network's training.

Is there consciousness without cognition? Is there cognition without consciousness? A behavioral approach would confirm this second possibility. A spiritualist approach would confirm the first. Most commonly people speak of cognition and consciousness as one, and think a clear consensus on whether they may be considered separately would greatly reduce the scope of contention when discussing the validity of artificial emulation of natural neural networks.

My objection to this insistence on a syntactic substrate for semantics is this: just because a model works doesn't mean it is correct or necessary. I may see a certain person as conforming to a certain stereotype, and behave as if this is true. That model may do well enough in most circumstances, even though the individuals concerned are nothing like the stereotype. I can characterize a number series as being the result of certain computations -- when in the particular case, they are due to some other calculation involving quite different criteria. If this number series happens to model some physical regularity, I may characterize it as a physical law, even though the actual factors are quite different. The model which includes syntax as a necessary substrate of semantics is such a construction. Syntax is an easy and convenient way to analyze semantic utterance, but it is only a model, not a constituent.

In a neural network, knowledge is not represented symbolically; nothing is acted on symbolically. The system response is holistic and, as Hopfield said, "Those neurons whose operation is dominantly linear merely provide a pathway of communication between nonlinear neurons."[136]

136. Hopfield. *Ibid.*

Appendix 4. Defining the Basic Computational Unit.

Shepherd's discussion of neuron versus synapse as the BCU:[137]
A common goal of experimental and theoretical neuroscience is to understand the neural basis of cognitive function. In pursuing this goal, experimentalists accumulating a wealth of information about the molecular and cellular properties of neurons, and about neural circuits and pathways in the brain. This information emphasizes the complexity of neurons, but gives only limited insight by itself into the way neurons form functional aggregates at the systems level to carry out cognitive functions. Network modelers, on the other hand, traditionally ignore most of the information gathered by experimentalists, replacing the elaborate structure of the neuron by a simple summing node, and the specific axonal and dendritic connections by a widely distributed network of wires. This provides a convenient approach to analyzing systems level behavior, but raises series doubts about its relevance to the real nervous system.

Shepherd's organizational levels:

A. Behavior
B. Systems and Pathways
C. Centers and local circuits
D. Neuron
E. Microcircuits
F. Synapse
G. Membranes, Molecules, Ions

The "functional unit" is the "morphological substrate for a specific function."

"The functional properties at each level are produced by the collective actions of the functional units derived from the next lower level."

Shepherd cites the collective action of muscle fibers as an example of emergent properties. i.e., properties not from individual units, but from collective action. (I don't think this is a good example of emergent properties -- see Appendix 2.[138])

"The relevance of a model for a real biological system depends on the extent to which its subcomponents are not arbitrary, but

137. Shepherd. *Ibid.*
138. Nagel, Ernest. *Ibid.*

represent properties that can be tested in the biological system."[139]
"If the organization at one level of the nervous system is thus constrained by its subcomponents at a lower level, it means that the properties at that level are not entirely emergent. For example, the force-generating properties of a whole muscle are implicit in the force-generating properties of its constituent muscle fibers. It becomes important . . . Therefore one cannot accurately simulate a given behavior . . . with a model system that does not incorporate the crucial levels of organization . . ."[140]
1) What are the basic computational units of nerve circuits
2) What operations are mediated by the different levels of organization
3) Which ones are critically relevant for modeling cognitive functions by the real nervous system
The basic computational unit of the nervous system is the synapse.

The Neuron Doctrine:
Traditionally the neuron is the basic anatomical, functional, metabolic, and developmental unit of the nervous system.
1) The cell body is all that matters as a representation of the integrative substrate of the neuron.
2) Axonal projections are arbitrary, depending on logic operations performed by the network.
3) Specific branching patterns of dendrites and axons of real neurons are simply ignored.
A lack of real neuron architecture implies little obvious relevance of these networks to real nervous systems. Rall (1959) claims the integrative properties of the neuron are dominated by the dendritic tree. Dendritic branches and dendritic spines are not simply passive . . . but are sites of dynamic biochemical changes of great relevance to memory and higher cognitive functions.

Neurons give off only one axon which has a definite pattern of collaterization within its local region and in its distant target regions. In many neurons the dendrites are the sites of synaptic outputs as well as inputs. . . Subcomponents of the dendrites can function as semi-independent I/O units. Functional units within neighboring dendrites can be linked to form local circuits involving only parts of each of the contributing neurons. These circuits may operate in the analog or digital or mixed modes. Therefore the Basic Computational Unit cannot be the neuron. The synapse is. This is the **Synaptic Doctrine**.

139. Shepherd. *Ibid.* 83.
140. Shepherd. *Ibid.* 84.

He concludes:

Computational circuits are built of synaptic circuits at successive levels of complexity; neurons thus represent one of the intermediate, multisynaptic levels.

Appendix 5. Learning.

Hebb cites James Mill with a theory of association a century prior to his own writing. The first object of Hebb's book was to present a theory of behavior for the consideration of psychologists; but equally important was his desire "to seek a common ground with the anatomist, physiologist, and neurologist, to show them how psychological theory relates to their problems and at the same time to make it more possible for them to contribute to the theory."[141]

> "Modern psychology takes completely for granted that behavior and neural function are perfectly correlated, that one is completely caused by the other. There is no separate soul or life-force to stick a finger into the brain now and then and make neural cells do what they would not otherwise. Actually, of course, this is a working assumption only -- as long as there are unexplained aspects of behavior." [142]

Hebb seems to have in mind the problem of emergent properties:

> "The problem of understanding behavior is the problem of understanding the total action of the nervous system, and *vice versa* . . . One must sympathize with those who want nothing of the psychologist's hair-splitting or the indefiniteness of psychological theory. There is much more certainty in the study of the electrical activity of a well-defined tract in the brain. The only question is whether a physiology of the brain as a whole can be achieved by such studies alone. One can discover the properties of its various parts more or less in isolation; but it is a truism by now that the part may have properties that are not evident in isolation, and these are to be discovered only by study of the whole intact brain."[143]

Hebb's theory is in part a reaction to two opposing formulations of brain function:

> "Two kinds of formula have been used, leading at two extremes to (1) switchboard theory, and sensori-motor connections; and (2) field theory . . . (1) In the first type of theory . . . the function of the cortex is that of a telephone exchange. Connections rigidly determine . . . what [a] human being does . . . (2) Theory at the opposite extreme denies that learning depends on connections at all . . . The cortex is regarded as made up of so many cells that it can

141. Hebb. *Ibid.* xi-xix.
142. Hebb. *Ibid.* xi-xix.
143. Hebb. *Ibid.* xi-xix.

be treated as a statistically homogeneous medium . . . both theoretical approaches seem to imply a prompt transmission of sensory excitation to the motor side . . . No one, at any rate, has made any serious attempt to elaborate ideas of a central neural mechanism to account for the delay, between stimulation and response, that seems so characteristic of thought."[144]

The principles of learning that Hebb expounded are as follows:

"Let us assume then that the persistence or repetition of a reverberatory activity (or "trace") tends to induce lasting cellular changes that add to its stability. The assumption can be precisely stated as follows: *When an axion (sic) of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.*"[145]

According to Kelso et al., a time-dependent, highly local, and strongly interactive mechanism is responsible for one form of *long-term potentiation* [LTP] in the hippocampus; there is strong neuropsychological evidence indicating that the *hippocampus* plays a key role in certain aspects of learning or memory.[146] Some studies support the theory that LTP in certain hippocampal synapses is Hebbian. Other studies show that the induction of synaptic LTP at some sites can be at variance with predictions based on Hebb's postulate of learning. Studies also show that the simple Hebbian learning rule emerges from a non-Hebbian LTP induction rule. Using computer simulations, Hetherington and Shapiro have demonstrated that

(1) an anti-Hebbian rule is needed to decrease the saturation of cell assembly activity,

(2) a synaptic modification rule that decreases synaptic weights when post-synaptic activity occurs in the absence of presynaptic activity is necessary, but not sufficient, for stable assemblies, and

(3) dendritic trees must be partitioned into independent regions of activation.

144. Hebb. *Ibid.* xi-xix.
145. Hebb. *Ibid.* 60-78.
146. Haykin. *Ibid.* 52-53.

Appendix 6. Computational Autonomy

According to Marr:[147]

> Almost never can a complex system of any kind be understood as a simple extrapolation from the properties of its elementary components. Consider, for example, some gas in a bottle. A description of thermodynamic effects -- temperature, pressure, density, and the relationships among these factors -- is not formulated by using a large set of equations, one for each of the particles involved. Such effects are described at their own level, that of an enormous collection of particles; the effort is to show that in principle the microscopic and macroscopic descriptions are consistent with one another. If one hopes to achieve a full understanding of a system as complicated as a nervous system, a developing embryo, a set of metabolic pathways, a bottle of gas, or even a large computer program, then one must be prepared to contemplate different kinds of explanation at different levels of description that are linked, at least in principle, into a cohesive whole, even if linking the levels in complete detail is impractical. For the specific case of a system that solves an information-processing problem, there are in addition the twin strands of process and representation, and both of these ideas need some discussion. A *representation* is a formal system of making explicit certain entities or types of information, together with a specification of how the system does this. And I shall call the result of using a representation to describe a given entity a *description* of the entity in that representation . . . the notion that one can capture some aspect of reality by making a description of it using a symbol and that to do so can be useful seems to me a fascinating and powerful idea . . . there is a trade-off; any particular representation makes certain information explicit at the expense of information that is pushed into the background any (sic) may be quite hard to recover.

Two different issues were confused in Marr's scheme:

1) As a *matter of discovery* can one figure out the algorithm and the problem analysis independently of facts about implementation? In fact, even if the *formal* algorithm is the same on two different architectures, the architectures themselves will affect factors of speed, size, efficiency, etc.

2) As a matter of formal theory, can a given algorithm that is already known to

147. Marr, David. *Vision*. San Francisco: W. H. Freeman, 1982. 19-38, 54-61.

perform a task in a given machine (e.g., the brain) be implemented in some other machine that has a distinct architecture?

Yes. By definition, a *formal* algorithm may be implemented on any sufficiently powerful computational architecture. It is quite possible to emulate symbolic programs using neural networks, and vice versa. Parallel and sequential programs are demonstrably interchangeable.

Bibliography

Alexander, R. McNeill. *Exploring Biomechanics: Animals in Motion.* New York: Scientific American Library. 1992.

Alexeieff, Alexander and Claire Parker. "Pinscreen Excerpts." (1973). *Canadian Animation Volume 2.* National Film Board of Canada: Pioneer LDCA, 1992.

Anderson, James A. and Edward Rosenfeld. *Neurocomputing: Foundations of Research.* Cambridge: The MIT Press. 1989.

Anderson, James A., Andras Pellionisz and Edward Rosenfeld. *Neurocomputing 2: Directions for Research.* Cambridge: The MIT Press. 1990.

Anderson, John R. *Cognitive Psychology and Its Implications.* New York: W. H. Freeman and Company. 1990.

Ang, Siew-Lan. "The brain organization." *Nature*, Vol 380, 7 March 1996. 25-27.

Angeles, Peter A. *Dictionary of Philosophy.* New York: Barnes and Noble Books. 1981.

Babel, Jean, Albert Bischoff and Heinrich Spoendlin. *Ultrastructure of the Peripheral Nervous System and Sense Organs.* Saint Louis, MO: The C. V. Mosby Company. 1970.

Bailey, Ronald H. *The Role of the Brain.* New York: Time-Life Books. 1975.

Baugh, Albert C. *A Literary History of England.* New York: Appleton-Century-Crofts, 1967.

Biehl, Soni. "Multiple Intelligences." TS. University of Texas at San Antonio.

Block, Ned. "On a confusion about a function of consciousness." *Behavioral and Brain Sciences* (1995) 18, 227-287.

Bourland, Hervé A. "How Connectionist Models Could Improve Markov Models for Speech Recognition", 1990 (in Eckmiller, *Advanced Neural Computers*, North-Holland, 1990.)

Brown, Douglas H. *Principles of Language Learning and Teaching, 2d ed.* Englewood Cliffs, NJ: Prentice Hall Regents, 1987, 68-70.

Caramazza, Alfonso. "The brain's dictionary." *Nature*, Vol 380, 11 April 1996. 485-486.

Causey, Robert L. *Logic, Sets, and Recursion.* Boston: Jones and Bartlett Publishers, 1994.

Churchland, Patricia, Christoph Koch and Terrence J. Sejnowski. "What Is Computational Neuroscience?" in Schwartz, Eric. *Computational Neuroscience.* Cambridge: the MIT Press. 1990.

Churchland, Patricia S. and Terrence J. Sejnowski. *The Computational Brain.* Cambridge: The MIT Press. 1992

Clarke, Stephanie, Anne Bellmann, François de Ribaupierre and Gil Assal. "Non-verbal auditory recognition in normal subjects and brain-damaged patients: Evidence for parallel processing." *Neuropsychologia*, Vol 34, No. 6, 587-603, 1996.

Cook, Norman D. *The Brain Code.* New York: Methuen, 1986.

Coulter, Jeff. *Rethinking Cognitive Theory.* New York: St. Martin's Press. 1983.

Crick, Francis. "Visual perception: rivalry and consciousness." *Nature*, Vol 379, 8 February 1996. 485-6.

Crick, Francis and Cristof Koch. "Some Reflections on Visual Awareness." *Cold Spring Harbor Symposia on Quantitative Biology*, Volume LV. 1990. 953-962.

---. "The Problem of Consciousness." *Scientific American*, Vol 267, Nr. 3, September 1992. 152-159.

---. "Towards a neurobiological theory of consciousness." *The Neurosciences*, Vol 2, 1990: 263-275.

Crosson, Frederick J. and Kenneth M. Sayre, eds. *Philosophy and Cybernetics.* New York: Simon and Schuster. 1967.

Damasio, Antonio R. and Hanna Damasio. "Brain and Language." *Scientific American*, Vol 267, Nr. 3, September 1992. 88-95.

Darden, Severn. *The Sound of My Own Voice and Other Noises.* Mercury LP OCS 6202. 1961.

Davidson, Richard J., Gary E. Schwartz and David Shapiro, eds. *Consciousness and Self-Regulation: Advances in Research and Theory*, Vol. 4 New York: Plenum Press. 1986.

Davis, Steven, ed. *Connectionism: Theory and Practice.* New York: Oxford University Press. 1992.

de Duve, Christian. *The Living Cell, Vol One.* New York: Scientific American Library. 1984.

De Mori, Renato. "Report on the paper 'Signals to Symbols to Meaning: Machine Understanding of Spoken Language' ", *Proceedings of the Congress of Phonetic Sciences*, Aix-en-Provence, France, AUG 19-24, 1991, Vol. 1. 84-86

Dennison, Paul E. & Gail E. Hargrove. *Personalized Whole Brain Integration.* Glendale, CA: Edu-Kinesthetics, Inc. 1985. 9.

Der, R. and M. Herrmann. "Q-Learning Chaos Controller." *IEEE*, 1994. 2472-2475.

Diamond, M. C., A. B. Scheibel, and L. M. Elson. *The Human Brain Coloring Book.* New York: Harper Perennial. 1985.

Dinsmore, John. *The Symbolic and Connectionist Paradigms: Closing the Gap.* Hillsdale, New Jersey: Lawrence Erlbaum Associates, Publishers 1992.

Dreyfus, Hubert L. "Phenomenology of Embodiment." *The Electronic Journal of Analytic Philosophy*, 4 (Spring 1996). Bloomington: Indiana University. 1996.

Drouin, Jacques. "Mindscape." (1976). *Canadian Animation Volume 2.* National Film Board of Canada: Pioneer LDCA, 1992.

Ellis, Andrew W. and Andrew W. Young. *Human Cognitive Neuropsychology.* Hillsdale: Lawrence Erlbaum Associates, Publishers, 1988.

Fallside, F. and W. A. Woods, eds., *Computer Speech Processing*, N.d.: Prentice Hall,

1985.
Fallside, Frank. "Progress in Large Vocabulary Speech Recognition", *Speech Technology*, APR/MAY 1989.
Fidone, Eyzaguirre. *Physiology of the Nervous System*. Chicago: Year Book Medical Publishers, Inc. 1969.
Fischbach, Gerald D. "Mind and Brain". *Scientific American*, September, 1992, 48-57.
Flanagan, Owen. *Consciousness Reconsidered*. Cambridge: MIT Press. 1992.
Franklin, Jon. *Molecules of the Mind*. New York: Athenaeum. 1987.
Fukushima, Kunihiko. "A Neural Network Model for Selective Attention." Tokyo: NHK Science and Technical Research Laboratories 1-10-11, Kinuta, Setagaya, Tokyo 157, Japan
Fuster, Joaquín M. *Memory in the Cerebral Cortex: An Empirical Approach to Neural Networks in the Human and Nonhuman Primate*. Cambridge: The MIT Press. 1995.
Garfield, Jay L., ed. *Foundations of Cognitive Science: The Essential Readings*. New York: Paragon House. 1990.
Gao, Jia-Hong, Lawrence M. Parsons, James M. Bower, Jinhu Xiong, Jinqi Li, Peter T. Fox. "Cerebellum Implicated in Sensory Acquisition and Discrimination Rather than Motor Control." *Science* VOL N.d. 1996, 1-3.
Gardner, Howard. *Frames of Mind: The Theory of Multiple Intelligences*. New York: Basic Books. 1983.
---. *The Mind's New Science: A History of the Cognitive Revolution*. New York: Basic Books. 1985.
---. *The Shattered Mind: The Person After Brain Damage*. New York: Vintage Books. 1976.
Gerefa. in Baugh, Albert C. *A Literary History of England*. New York: Appleton-Century-Crofts, 1967, 36.
Gershon, Elliot S. and Ronald O. Rieder. "Major Disorders of Mind and Brain." *Scientific American*, Vol 267, Nr. 3, September 1992. 126-133.
Gleick, James. *Chaos: Making a New Science*. New York: Penguin Books. 1988.
Goldman-Rakic, Patricia S. "Working Memory and the Mind." *Scientific American*, Vol 267, Nr. 3, September 1992. 110-117.
Gray, Henry. *Anatomy*. New York: Bounty Books. 1977.
Guthrie, Louise, James Pustejovsky, Yorick Wilks, and Brian M. Slator. "The Role of Lexicons in Natural Language Processing." *Communications of the ACM*. JAN 1996, 63-72.
Harnad, Stevan. "Grounding Symbols in the Analog World with Neural Nets: a Hybrid Model." *Think Quarterly*, Vol. 2, No. 1, 1993. Instituut voor Taal- en Kennistechnologie (http://tkiwww.kub.nl:2080/tki/Docs/Think)
Haugeland, John, ed. *Mind Design: Philosophy, Psychology, Artificial Intelligence*.

Cambridge: The MIT Press. 1988.

Haykin, Simon. *Neural Networks - A Comprehensive Foundation*. New York: Macmillan College Publishing Company, 1994.

Hebb, Donald O. *The Organization of Behavior*. New York: Wiley, Introduction and Chapter 4, "The first stage of perception: growth and assembly," 1949. xi-xix, 60-78.

Hilgard, Ernest R. *Theories of Learning*. New York: Appleton-Century-Crofts, Inc. 1956.

Hinton, Geoffrey E. "How Neural Networks Learn from Experience." *Scientific American*, Vol 267, Nr. 3, September 1992. 144-151.

Hofstadter, Douglas. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York: Basic Books. 1995.

---. *Gödel, Escher, Bach: an Eternal Golden Braid*. New York: Vintage Books. 1980.

---. *What is the Meaning of the Word 'I'?* The Grace A. Tanner Lecture in Human Values. Cedar City, Utah: Southern Utah State College. 1982. III.

---. *Metamagical Themas: Questing for the Essence of Mind and Pattern*. New York: Basic Books, Inc., Publishers. 1985.

---. *The Mind's I: Fantasies and Reflections on Self and Soul*. New York: Bantam Books. 1982.

Honavar, Vasant and Leonard Uhr. "Brain-structured connectionist networks that perceive and learn." Computer Sciences Technical Report # 843. Madison: University of Wisconsin. April 1989.

---. "Experimental Results Indicate that Generation, Local Receptive Fields and Global convergence Improve Perceptual Learning in Connectionist Networks." Computer Sciences Technical Report # 805. Madison: University of Wisconsin. November 1988.

---. "A Note on the Symbol Grounding Problem and its Solution." *Think Quarterly*, Vol. 2, No. 1, 1993. Instituut voor Taal- en Kennistechnologie (http://tkiwww.kub.nl:2080/tki/Docs/Think)

---. "A Network of Neuron-Like Units that Learns to Perceive by Generation as Well as Reweighing of its Links." Computer Sciences Technical Report # 793. Madison: University of Wisconsin. September 1988.

Hopfield, J. J. "Neural networks and physical systems with emergent collective computational abilities." *Proceedings of the National Academy of Sciences* 79: 2554-2558. 1982.

Huang, X. D., Y. Ariki, and M. A. Jack, *Hidden Markov Models for Speech Recognition*, Edinburgh University Press, 1990.

Huang, Zezhen and Anthony Kuh, "A Combined Self-Organizing Feature Map and Multilayer Perceptron for Isolated Word Recognition", *IEEE Transactions On Signal Processing*, VOL. 40, NO. 11, November, 1992.

Hubel, David H. "A big step along the visual pathway." *Nature.* Vol 380, 21 March 1996. 197-198.

---. *Eye, Brain, and Vision.* New York: Scientific American Library. 1988.

Jacquette, Dale. *Philosophy of Mind.* Englewood Cliffs, NJ: Prentice Hall, 1994.

Johnstone, Rufus A. "Female preference for symmetrical males as a by-product of selection for mate recognition." Nature, VOL 372, 10 NOV 1994. 172-175.

Kandel, Eric R. and Robert D. Hawkins. "The Biological Basis of Learning and Individuality." *Scientific American,* Vol 267, Nr. 3, September 1992. 78-86.

Kangas, Jari. *On the Analysis of Pattern Sequences by Self-Organizing Maps* (Doctoral Thesis), Helsinki University of Technology, 1994

Kangas, Jari; Kari Torkkola and Mikko Kokkonen, "Using SOMS as Feature Extractors for Speech Recognition". N.d.

Kasabov, N. and E. Peev. "Phoneme Recognition with Hierarchical Self Organised Neural Networks and Fuzzy Systems - A Case Study", in *ICANN '94, Proceedings of the International Conference on Artificial Neural Networks,* Vol 1, Part 2., pp. 201 - 204

Kimura, Doreen. "Sex Differences in the Brain." *Scientific American,* Vol 267, Nr. 3, September 1992. 118-125.

King, P. J. and E. H. Mamdani. "The Application of Fuzzy Control Systems to Industrial Processes." NT, SEP 1976. 235-242.

Klatt, D. H. "Review of Text-To-Speech Conversion for English", *JASA* 82(3):737-793, Sept 1987.

Klimesch, Wolfgang. *The Structure of Long-Term Memory.* Hillsdale: Lawrence Erlbaum Associates, 1994, 12.

Knecht, S., E. Kunesch and A. Schnitzler. "Parallel and serial processing of haptic information in man: Effects of parietal lesions on sensorimotor hand function." *Neuropsychologia,* Vol. 34, No. 7, 669-687, 1996.

Knowles, Dudley, ed. *Explanation and its Limits.* Cambridge: Cambridge University Press. 1990.

Kohonen, Teuvo. "The Adaptive-Subspace SOM (ASSOM) and its Use for the Implementation of Invariant Feature Detection." Paris: *Proc. ICANN* '95. 1995.

Kong, Seong-Gon and Bart Kosko. "Comparison of Fuzzy and Neural Truck Backer-Upper Control Systems." Ch. 9. n.d. 339-361.

Kosslyn, Stephen M. and Olivier Koenig. *Wet Mind: The New Cognitive Neuroscience.* New York: The Free Press. 1995.

Leahey, Thomas H. and Richard Jackson Harris. *Learning and Cognition.* Englewood Cliffs, NJ: Prentice Hall. 1993.

Lee, K-F. *Automatic Speech Recognition: The Development of the SPHINX System,* Kluwer, 1989.

Lippman, Richard P. "Review of Neural Networks for Speech Recognition", *Neural Computation* 1(1): 1-38, 1989, MIT.

Lynch, Gary. *Synapses, Circuits, and the Beginnings of Memory.* Cambridge: The MIT Press. 1986.

Mack, Walter N. and Elizabeth A. Leistikow. "Sands of the World." *Scientific American.* Volume 275, Number 2, August 1996

Mammone, Richard J., Ed. *Artificial Neural Networks of Speech and Vision,* Chapman & Hall, London. 1994.

Marr, David. *Vision.* San Francisco: W. H. Freeman, 1982. 19-38, 54-61.

Martini, Frederick H. *Fundamentals of Anatomy and Physiology, 3d ed.* Englewood Cliffs: Prentice Hall. 1995, 468.

McCarthy, Rosaleen A. and Elizabeth K. Warrington. *Cognitive Neuropsychology: A Clinical Introduction.* San Diego: Academic Press. 1990.

McClelland, James L. and Jeffrey L. Elman. "The TRACE Model of Speech Perception", *Cognitive Psychology* 18, 1-86, 1986.

Merleau-Ponty, Maurice. *The Primacy of Perception.* Northwestern University Press. 1964.

Minsky, Marvin. *The Society of the Mind.* New York: Simon and Schuster. 1986.

Miikkulainen, Risto. *DISCERN: A Distributed Artificial Neural Network Model of Script Processing and Memory.* PhD Dissertation, Technical Report UCLA-AI-90-05. Los Angeles: Computer Science Department, University of California. September 1990.

---. "A Distributed Feature Map Model of the Lexicon." *Proceedings of the 12th Annual Conference of the Cognitive Science Society*, 1990.

---. "Subsymbolic Case-Role Analysis of Sentences with Embedded Clauses." *Cognitive Science*, in press.

---. *Subsymbolic Natural Language Processing: An Integrated Model of Scripts, Lexicon, and Memory.* Cambridge: MIT Press, 1993.

Miikkulainen, Risto and Michael G. Dyer. "Encoding Input/Output Representations in Connectionist Cognitive Systems." Technical Report UCLA-AI-88-16. Los Angeles: Artificial Intelligence Laboratory, Computer Science Department, University of California. 1988.

---. "Natural Language Processing With Modular PDP Networks and Distributed Lexicon." *Cognitive Science*, 15 (3): 343-399, 1991.

Miikkulainen, Risto and Gert Westermann. "Verb Inflections in German Child Language: A Connectionist Account." *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, 1994.

Millikan, Clark H., Chairman and Frederic L. Darley, ed. *Brain Mechanisms Underlying Speech and Language.* New York: Grune and Stratton. 1967.

Moriarty, David E. and Risto Miikkulainen. "Efficient Reinforcement Learning through Symbiotic Evolution." *Machine Learning* 22, 1996. Boston: Kluwer Academic Publishers. 1996.

Mukerjee, Madhusree. *Scientific American,* Vol 275, Nr. 5, November 1996. 28-36.

Nagel, Ernest. *The Structure of Science*. New York: Harcourt, Brace & World, Inc. 1961.

Nix, David A. and Andreas S. Weigand. "Estimating the Mean and Variance of the Target Probability Distribution." *IEEE*. 1994.

O'Connor, D. J. and Brian Carr. *Introduction to the Theory of Knowledge*. Brighton, Sussex: The Harvester Press. 1982.

Ornstein, Robert and Richard F. Thompson. *The Amazing Brain*. Boston: Houghton Mifflin. 1984.

O'Shaughnessy, D. *Speech Communication: Human and Machine*, Addison-Wesley, MA, 1987.

Palm, Günther and Ad Aertsen. *Brain Theory*. New York: Springer-Verlag. 1986.

Penfield, Wilder. *The Mystery of the Mind*. Princeton: Princeton University Press, 1975.

Penfield, Wilder and Lamar Roberts. *Speech and Brain-Mechanisms*. Princeton: Princeton University Press. 1959.

Phillips, John L., Jr. *The Origins of Intellect: Piaget's Theory*. San Francisco: W. H. freeman and Company. 1975.

---. *Piaget's Theory: A Primer*. San Francisco: W. H. Freeman and Company, 1981.

Pinker, Stephen. *The Language Instinct: How the Mind Creates Language*. New York: Harper Perennial, 1995.

Pinker, Stephen and Jacques Mehler, eds. *Connections and Symbols*. Cambridge: MIT Press, 1988.

Pons, Tim. "Novel sensations in the congenitally blind." *Nature*, Vol 380, 11 April 1996. 479-480.

Posner, M. I. & M. E. Raichle. "Précis of *Images of Mind.*" *Behavioral and Brain Sciences* (1995) 18, 327-383.

Quinlan, Philip T. *Connectionism and Psychology: A Psychological Perspective on New Connectionist Research*. Chicago: University of Chicago Press. 1991.

Rabiner, L. R. and R. W. Schafer. *Digital Processing of Speech Signals*, Prentice Hall, 1993.

Rabiner, L. R. and B-H Juang. *Fundamentals of Speech Recognition*, Prentice Hall, 1993.

Rahmann, H. and M. Rahmann. *The Neurobiological Basis of Memory and Behavior*. New York: Springer-Verlag. 1992, 220.

Rapaport, David. *Emotions and Memory*. New York: International Universities Press. 1971.

Renals, Steve and Nelson Morgan. "Connectionist Probability Estimation in HMM Speech Recognition", *ICSI*, TR-92-081, December 1992.

Restak, Richard M. *The Brain: The Last Frontier*. Garden City, New York: Doubleday & Company, Inc. 1979.

---. *Brainscapes*. New York: Hyperion. 1995.

---. *The Mind.* New York: Bantam Books. 1988.
---. *The Modular Brain.* New York: Touchstone. 1994.
Revonsuo, Antti, Matti Kamppinen and Seppo Sajama. *Consciousness in Philosophy and Cognitive Neuroscience.* Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers. 1994.
Rich, Elaine and Kevin Knight. *Artificial Intelligence.* New York: Mcgraw-Hill, Inc. 1991.
Rock, Irvin. *Perception.* New York: Scientific American Library. 1984.
Rumelhart, David E. *Towards a Microstructural Account of Human Reasoning.* Ch. 3 in Davis, Steven. *Connectionism: Theory and Practice.* Oxford University Press, New York, 1992.
Ryle, Gilbert. *The Concept of Mind.* Chicago: The University of Chicago Press. 1949.
Scarborough, Don L., Ben O. Miller, and Jacqueline A. Jones. "Connectionist Models for Tonal Analysis." in Todd et al (q.v.) *Music and Connectionism.* 1989.
Schafer, J. David, Darrell Whitely, and Larry J. Eshelman. *Combinations of Genetic Algorithms and Neural Networks: A Survey of the State of the Art.* Los Alamitos, CA: IEEE Computer Society Press. 1992.
Schafer, R. W., & J. D. Markel, eds., *Speech Analysis,* IEEE Press, 1979.
Schwartz, Eric L., ed. *Computational Neuroscience.* Cambridge: MIT Press. 1990.
Searle, John. *Minds, Brains and Science.* Cambridge: Harvard University Press. 1984.
---. "The Failures of Computationalism." (Response to Stevan Harnad (q.v.), "Grounding Symbols . . .") *Think Quarterly,* Vol. 2, No. 1, 1993. Instituut voor Taal- en Kennistechnologie (http://tkiwww.kub.nl:2080/tki/Docs/Think)
Seidenberg, Mark S. *Connectionism without Tears.* Ch. 4 in Davis, Steven. *Connectionism: Theory and Practice.* Oxford University Press, New York, 1992.
Siegelman, Hava. "The Super-Turing Result." *Science,* n.d. 1994.
Selkoe, Dennis J. "Aging Brain, Aging Mind." *Scientific American,* Vol 267, Nr. 3, September 1992. 134-142.
Serebriakoff, Victor. *Brain.* London: Davis-Poynter. 1975.
Shatz, Carla J. "The Developing Brain." *Scientific American,* Vol 267, Nr. 3, September 1992. 60-67
Shepherd, Gordon M. *Neurobiology.* New York: Oxford University Press. 1988.
---. "The Significance of Real Neuron Architecture for Neural Network Simulations." in Eric Schwartz (q.v.) *Computational Neuroscience.* 82-96
Shinbrot, Troy, Celso Grebogi, Edward Ott and James A Yorke. "Using small perturbations to control chaos." *Nature,* Vol 363, 3 JUNE 1993. 411-417.
Shrobe, Howard E. *Exploring Artificial Intelligence.* San Mateo, CA: Morgan Kaufman Publishers, Inc. 1988.
Smolensky, Paul. "On the proper treatment of Connectionism." *Behavioral and*

Brain Sciences (1988) 11, 1-74.
Soler, O. & R. van Hoe. "BAR: A Connectionist Model of Bilingual Access Representations", in *ICANN '94, Proceedings of the International Conference on Artificial Neural Networks*, Vol 1, Part 2. 263-267
Springer, Sally P. and Georg Deutsch. *Left Brain, Right Brain*. New York: W. H. Freeman and Company. 1989.
Swerdlow, Joel L. "Quiet Miracles of the Brain." *National Geographic* Vol. 187, No. 6. June 1995. 2-41.
Tanney, Rick W., "On the Relation of Ego Boundary and Neurosis." Master's Thesis. Ashland Theological Seminary, 1974.
---. "Quirks and Momentos: a Discontinuous Universe." In manuscript.
---. "Toward a Connectionist Lexicon and the Possibility of Disambiguation." In Manuscript.
Todd, Peter M. and D. Gareth Loy. *Music and Connectionism*. Cambridge: The MIT Press. 1991.
Totora, Gerard J. *Principles of Human Anatomy, 6th ed.* New York: Harper Collins Publishers Inc. 1992.
Umbanhowar, P. B., F. Melo & H. L. Swinney. "Localized excitations in a vertically vibrated granular layer." *Nature* 382, 793-796. 1996.
Vander, Arthur J., James H. Sherman and Dorothy S. Luciano. *Human Physiology: The Mechanisms of Body Function*. New York. McGraw-Hill Book Company. 1980.
von der Malsburg, Christoph. "Sensory segmentation with coupled neural oscillators." *Biological Cybernetics* 67, 233-242 (1992) Springer-Verlag 1992.
Walter, W. Grey. *The Living Brain*. New York: W. W. Norton & Company Inc. 1963.
Waibel, A. *Prosody and Speech Recognition*, Morgan Kaufman, 1988.
Waibel, Alex. "Modular Construction of Time-Delay Neural Networks for Speech Recognition", *Neural Computation* 1, 39-46, 1989 MIT.
Waibel, Alexander, Toshiyuki Hanazawa, Geoffrey Hinton, Kiyohiro Shikano, and Kevin J. Lang , "Phoneme Recognition Using Time-Delay Neural Networks", *IEEE Transactions on Acoustics, Speech, and Signal Processing*, VOL. 37, NO. 3, 1989.
Waibel, A. and K-F. Lee, eds., *Readings in speech Recognition*, Morgan Kaufman, 1990.
Waibel, Alex; Louise Osterholtz, Hiroaki Saito, Otto Schmidbauer, Tilo Sloboda, and Monika Woszczyna. "JANUS: Speech-to-Speech Translation Using Connectionist and Non-Connectionist Techniques"
Wasserman, n.n. *Neural Computing: Theory and Practice*. n.p. van Nostrand Reinhold. 1989. 127-149.
Wermter, Stefan. "Integration of Semantic and Constraints for Structural Noun

Phrase Disambiguation." *Speech and Neural Language* 1486-1491. nd.
West, Thomas G. *In the Mind's Eye*. Buffalo, New York: Prometheus Books. 1991.
Wolf, Fred Alan. *The Dreaming Universe*. New York: Touchstone/Simon & Schuster. 1994.
Wright, J. J., and D. T. J. Liley. "Dynamics of the brain at global and microscopic scales: Neural networks and the EEG." *Behavioral and Brain Sciences* (1996) 19, 285-320.
Zeki, Semir. "The Visual Image in Mind and Brain." *Scientific American*, Vol 267, Nr. 3, September 1992. 68-76.
Zue, Victor W. "From Signals to Symbols to Meaning: On Machine Understanding of Spoken Language", *Proceedings of the XIIth Congress of Phonetic Sciences*, Aix-en-Provence, France, AUG 19-24, 1991, Vol. 1 pp. 74-86
na. "A Concurrent Hybrid Genetic/Neural Network Learning algorithm for MIMD Shared Memory Machines." *Machine Learning*. 127-143. nd.
WWW: web sites with pertinent information or links:
ftp://ftp.ira.uka.de/pub/neuron/
http://neuro.med.cornell.edu/VL/
http://tkiwww.kub.nl:2080/tki/Docs/Think
http://www.comlab.ox.ac.uk/archive/
http://www.cs.utexas.edu/usrs/ailab
http://www.cs.reading.ac.uk/people/dwc/ai.html
http://www.cs.utexas.edu/users/nn
http://www.cs.utoronto.ca/~zoubin/
http://www.sei.cmu.edu/
http://www.yahoo.com/Science/Engineering/Electrical_Engineering/Neural Networks/

VITA

Rick Willard Tanney was born in New Castle, Pennsylvania on May 28, 1948, the son of Lillian Rose Fullwood and Alan Ellsworth Porter. He changed his name from Alan Ellsworth Porter, Junior, to include his step-father's (Willard Robert Tanney) first and last names. After completing his work at Lexington High School, Lexington, Ohio, in 1966, he entered the Mansfield Branch of The Ohio State University in Mansfield, Ohio, continuing after two years at the main campus in Columbus, Ohio, where he received a Bachelor of Arts in English in June, 1970. While working as a Psychological Nurse Technician at the Mansfield General Hospital, he entered the graduate school of Ashland College, Ashland Theological Seminary, in Ashland, Ohio where he received a Master of Divinity in Pastoral Counseling and Psychology in June, 1974. His thesis was entitled "On the Relation of Ego Boundary and Neurosis." He enlisted in the United States Army in September, 1974, and was stationed at the US Army Research Institute of Environmental Medicine in Natick, Massachusetts. While there, he participated in many studies of environmental injury and stress relating to cold weather, high altitude, work stress, and jet lag. By invitation of the local Arts Workshop, he held poetry workshops, and self-published a collection of his own work entitled "Re-Verse #1" in 1977. In April, 1981, he transferred to the 2d General Hospital, Landstuhl Regional Army Medical Center in Landstuhl, West Germany, where he was the Non-Commissioned Officer in Charge of the Mental Health Clinic. While there, he entered the graduate program of Boston University, Metropolitan College, and received a Master of Science in Computer Information Systems in 1984. In April, 1984, he transferred to Ireland Army Hospital at Fort Knox, Kentucky, where he was the Non-Commissioned Officer in Charge of the Social Work Service. In March, 1985, he transferred to the US Army Military Community Activity in Nürnburg, West Germany, where he successively held the jobs of Non-Commissioned Officer in Charge of the Alcohol and Drug Abuse Prevention Control Program, Education Coordinator, and Assistant Education Coordinator. He began teaching Computer Science for the University of Maryland, European Division. In September, 1987, he transferred to the US Army Research Unit-Europe in Heidelberg, West Germany, where he was Detachment Sergeant. He supervised research assistants in studies involving grief recovery after training accidents and various statistical analyses. He entered the undergraduate program of the University of Maryland, and received a Bachelor of Science degree in Computer Science in 1990, while continuing to teach for the University in Heidelberg and at remote locations in Germany. In November, 1990, he transferred to Brooke Army Medical Center at Fort Sam Houston, Texas, where he was successively Administrative Non-Commissioned Officer for Social Work Service, Non-Commissioned Officer in Charge of the Mental Health Clinic, Special Projects Non-Commissioned Officer and Non-commissioned Officer in Charge for the Department of Psychiatry, and Assistant Administrative Non-Commissioned Officer for Headquarters, Brooke Army Medical Center. In Fall 1992, he enrolled in The Graduate School at the University of Texas in Austin. In September , 1993, he transferred to the 85th Combat Stress Control Detachment at Fort Hood, Texas, where he served as Detachment First Sergeant for the activation of the unit. In November, 1993, he was appointed Information Management Officer for the 61st Area Support Medical Battalion at Fort Hood. He retired from the US Army with twenty years of active duty service with the rank of Sergeant First Class on the First of October, 1994. He teaches Computer Science for Palo Alto Community College in San Antonio, Texas and freelances as a computer consultant and instructor in the San Antonio area.

Permanent Address: 8111 Zigler Road
Sterling, Ohio 44276-0073
rtanne01@interserv.com

This report was word-processed by the author using a 90-MHz Pentium platform running Wordstar for Windows, and was printed on an Okimate OL400.