REPRINTED FROM

Elsevier's

Encyclopedia of Neuroscience

2nd enlarged and revised edition

Edited by
George Adelman
Barry H. Smith

Editorial Manager
Jennifer De Pasquale

Visit the Encyclopedia's website for supplementary information and details regarding the CD-ROM version:

http://www.elsevier.com/locate/encneu

Table 1. Major structural-biochemical causes of stupor and coma

Supratentorial mass lesions (secondarily producing deep diencephalic and/or upper brainstem dysfunction)
Cerebral hemorrhage
Large cerebral infarction
Traumatic brain contusion-edema
Subdural hematoma
Epidural hematoma
Brain tumor
Brain abscess (rare)
Subtentorial lesions (compressing or destroying ascending activating systems)
Pontine or cerebellar hemorrhage
Rostral pontine tegmental brain stem infarction
Brain stem or cerebellar expanding tumor
Cerebellar abscess compressing brain stem
Metabolic and diffuse lesions
Global cerebral anoxia or ischemia (e.g., cardiac arrest)
Hypoglycemia
Severe nutritional deficiency (e.g., advanced Wernicke's disease)
Endogenous organ failure or deficiency (e.g., lung, liver, kidney)
Exogenous poison (e.g., alcohol, sedatives, oplates)
Infections
Meningitis
Encephalitis
Ionic and electrolyte disorders (e.g., hyponatremia, water intoxication)
Status epilepticus
Concussion and postictal states

jury do better than those older than 30–40 years of age. By contrast, even in children, severe anoxic coma often leads to a permanently crippled existence. Limited degrees of structural damage such as occur with traumatic concussion can be followed relatively promptly by a complete or near complete return of brain function. More severe injury, especially when it produces structural damage to the cerebrum or critical brain stem regions usually results in long lasting cognitive impairment, the most severe example of which is the cognitively and emotionally empty condition called the vegetative state.

2. Etiology and management

Coma as a medical problem implies the imminent threat of brain failure with widespread loss of cerebral activity, upper brain stem function, or both. Table 1 lists the major diseases or categories that can produce such severe impairments. Supratentorial mass lesions per se generally impair consciousness little or not at all unless they expand and distort the brain sufficiently to compress the diencephalon latero-caudally, thereby producing transtentorial distortion or herniation. As one would anticipate, rapidly enlarging lesions are more dangerous in this respect than are relatively slowly changing ones. Upper brain stem or lower diencephalic abnormalities cause coma when the lesions directly damage the central ascending forebrain activating systems. Metabolic disorders can affect both the supra- and subtentorial mechanisms that normally generate conscious behavior. Accordingly, they generally produce mutifocal symptoms and signs, reflecting dysfunction at several anatomic levels of the brain. When briefly lasting, they may result in no sustained cerebral insufficiencies. Metabolic suppression of the brain by drugs or surgical anesthesia, for example, leaves in its wake no discoverable ill effects.

Most brain dysfunction sufficient to produce coma implies a poor prognosis; only therapeutic anesthesia or overdose with sedatives or alcohol contradict the rule. Among large series of patients with nontraumatic coma lasting more than 12 hours or so, only about 15% completely recovered their physical and intellectual functions. Likewise, among patients with severe, sustained coma from head injury, almost half will die and as many as 25% of the survivors will be severely incapacitated. Patients with metabolic coma or young persons briefly unconscious from head trauma generally do best, while those showing signs of severe primary or secondary brain stem damage fare worse, no matter what the cause.

Patients in coma are best treated acutely in special care units where they can receive close attention to their often precarious autonomic functions and can obtain specific therapy directed at their underlying neurological disease. Skilled respiratory and cardiovascular support provide the necessary core of management with specific measures directed at the particular disease that threatens to destroy the brain.

Further reading

Plum F. (1991): Coma and related global disturbances of the human conscious state. In: *Cerebral Cortex, Vol. 9,* Peters, A, ed.New York: Plenum

Plum F, Posner JB (1997): *The Diagnosis of Stupor and Coma.* 4th edn. Philadelphia: FA Davis

See also Persistent vegetative state; Activation, arousal, alertness and attention; Brain death; Brain trauma; Hypoxia; Brain injury, functional recovery after

Computational neuroscience

Terrence J. Sejnowski

Computational neuroscience is a relatively recent approach to understanding how nervous systems represent, process, store, and act upon information that is latent in the environment or is expressed genetically through developmental mechanisms. Models of neural systems can be used to interpret experimental data in new ways, to confirm and extend existing hypotheses, and to generate new hypotheses for the function of neural systems. These hypotheses provide links between levels of description, from the molecular to systems levels. The ultimate aim of computational neuroscience is to provide linking principles from neural mechanisms to behavior.

Table 1. Resources for computational neuroscience

Annual Summer Schools and Conferences
Summer Course on Methods in Computational Neuroscience; Woods Hole MA (August)
Crete Course on Computational Neuroscience; Crete, Greece (September)
Cold Spring Harbor Laboratories Summer Course on Computational Neuroscience: Vision; Cold Spring Harbor, New York (July)
Neural Information Processing Systems Conference; Denver Colorado (November)
Computation and Neural Systems Conference, Alternates between East Coast and West Coast (July)
Selected Journals
Journal of Computational Neuroscience (Kluwer Academic Publishers)
Network: Computation and Neural Systems (IOP Press)
Neural Computation (MIT Press)
Simulation Programs
MCell (Thomas Bartol and Joel Stiles): Monte Carlo models of subcellular chemical signaling
GENESIS (James Bower and Matthew Wilson): Realistic compartmental models of neurons and networks
NEURON (Michael Hines and John Moore): Realistic compartmental models of neurons and networks
NSL (Michael Arbib and Alfredo Weitzenfeld): Neural Simulation Language for large-scale models of neural systems
PDP++ (Randall O'Reilly and James McClelland): Parallel Distributed Processing models based on abstract neural networks

Computational neuroscience has made progress in achieving these aims by using techniques from computer science and applied mathematics to simulate and analyze computational models of neurons and neural systems at many levels of investigation. Digital computers have continued to increase in speed, making it possible to approach more complex neural systems. The number of investigators using computational tools is expanding and a variety of new journals, summer schools and scientific conferences have proliferated that focus on computational neuroscience (Table 1). A comprehensive handbook on brain theory has appeared (Arbib, 1995). In this article, only a few of the major issues and advances in the field can be summarized.

1. What is computation?

Many different types of physical systems can solve computational problems, including slide rules and optical Fourier analyzers as well as digital computers. What these have in common is an underlying correspondence between an abstract mathematical algorithm and the states of the physical system (Churchland and Sejnowski, 1992). This approach to computation is broad enough to include neural systems. An important distinction can be made between general purpose computers, which can be programmed to solve many different types of algorithms, and special purpose computers, which are designed to solve only a limited range of problems. Most neural systems are specialized for particular tasks, such as the retina which is dedicated to visual transduction and image processing. Because of the close coupling between structure and function in a dedicated system, the anatomy and physiology of a brain region provide important clues to its function. Unlike a digital computer, the connectivity between neurons and their properties are shaped by the environment during development and remain plastic even in adulthood. Thus, as the brain processes information, it changes its own structure in response to the information. This plasticity is important in allowing brains to respond flexibly to a changing world through adaptation and learning.

2. Brain modeling

Brain models used as an adjunct to experimental techniques have several advantages: (1) Models provide intuition about the possible behaviors of complex, dynamical brain systems, especially when they are nonlinear and have feedback loops; (2) the predictions of a model make explicit the consequences of the underlying assumptions, and comparison with experimental results can lead to new insights and discoveries; and (3) the results of difficult experiments can be simulated with a model, such as reversible lesions of selected channels or neurons, to optimize the design of the experiment for distinguishing between competing explanations.

Eve Marder has made an interesting distinction between three different types of brain models. The first type, called an interpretive model, is used to analyze experimental data in order to determine whether they are consistent with a particular computational assumption. For example, Apostolos Georgopoulos has used a "vector averaging" technique to compute the direction of arm motion from the responses of a population of cortical neurons, and William Newsome and his colleagues have used signal detection theory to analyze the information from single cortical neurons responding to visual motion stimuli. In these examples, the computational model was used to explore the information in the data but was not meant to be a model for the actual cortical mechanisms. Nonetheless, these models were highly influential and have provided new ideas for how the cortex may represent sensory information and motor commands. These models in turn have affected experimental design, which have then led to improved models.

A second type of model, called a confirmatory model, has been used extensively to test whether a set of data can account for the phenomena being studied. In many biophysical experiments, such as the classic Hodgkin-Huxley studies of the squid action potential, sets of data are collected under a variety of conditions, and a model is later constructed to integrate the data into a unified framework. This type of model is most effective when most of the variables in the model have been measured experimentally and only a few unknown parameters need to be fit to the experimental data. One danger with this approach is that even if the model fits the data, the resulting model may not be unique. However, automated techniques have been developed for systematically exploring large parameter spaces to determine all combinations of parameters that fit the data (Koch and Segev, 1997). As the number of experiments increases, the number of possible solutions that fit all the data should converge to a unique set.

Finally, a third type of model starts with a general principle and produces a speculative model that implements the principle within known biological constraints. These models can be quite fruitful in helping to motivate experiments that might not have been otherwise undertaken. An example of this approach is the model of coupled nonlinear oscillators analyzed by Nancy Kopell and others, which has led to new experiments on fictive swimming in the lamprey spinal cord (Churchland and Sejnowski, 1992). One of the strengths of this type of model is that it can be used to identify the critical variables that determine the qualitative behavior of a system. One corresponding weaknesses is that because much of the fine detail is often absent, it may not be possible to make detailed comparisons with data.

Dynamical systems theory has been used to analyze dynamical models of small neural systems. Dynamical systems analysis is most fruitful when the numbers of parameters and variables are small. Most models of neural networks involve a large number of variables, such as membrane potentials, firing rates, and concentrations of ions, with an even greater number of unknown parameters such as synaptic strengths, time constants, and conductances. In the limit that the number of neurons and parameters is very large, techniques from physics become applicable in predicting the average behavior of large networks. There is a midrange of systems where neither type of limiting analysis is possible, but where simulations can be performed. One danger of relying solely on computer simulations is that the they may be as complex and difficult to interpret as the biological system itself.

3. Recent brain models

At the cellular and molecular levels, significant advances have taken place in modeling specific neurons and synapses based on biophysical measurements of ionic mechanisms. Markov models that are used to model ionic channels can be applied to every aspect of synaptic signaling, including transmitter release and the intracellular second-messenger systems that modulate synaptic transmission (Koch and Segev, 1997). The original Hodgkin-Huxley models for the fast sodium and delayed rectifier potassium channels are special cases of a Markov model, as are the detailed biophysical models of receptor kinetics.

The pioneering work of Wilfrid Rall on the electrical properties of dendrites was based on the analysis of simplified dendritic geometries (Segev et al., 1995). It is now possible to simulate multicompartment models of dendrites from the geometries of reconstructed neurons. Voltage-dependent sodium and calcium channels have been observed in the dendrites of cortical neurons, which greatly increases the complexity of synaptic integration. The experimental finding that active currents can carry information in a retrograde direction from the cell body up to the distal synapses also has computational significance for Hebbian forms of synaptic plasticity. Another intriguing observation made with modeling techniques is that the wide variety of spiking patterns in cortical neurons can be reproduced from the same distribution of ionic channels by varying only the geometry of the dendritic tree (Koch and Segev, 1997).

Realistic models with several thousand cortical neurons can be explored on the current generation of workstation, which allows the dynamics of cortical columns can be explored in detail. The first model for the orientation specificity of neurons in the visual cortex was the feedforward model proposed by Hubel and Wiesel, which assumed that the orientation preference of cortical cells was determined primarily by converging inputs from thalamic relay neurons. Experimental evidence now favors this model over other models in which the orientation specificity was determined primarily by local cortical circuits. Simulations of orientation columns have shown that the intrinsic circuits in the cortex could be used to amplify weak signals and suppress noise as well as to perform gain control to extend its dynamic range. This is an important step forward in understanding the function of visual cortex. Abstract models, such as Hopfield networks, have also been helpful in providing a conceptual framework for cortical dynamics (Churchland and Sejnowski, 1992).

Although thalamic neurons that project to the cortex are called relay cells, they almost surely have additional functions since the visual cortex makes massive feedback projections back to them. As an example of a speculative model that has led to a new computational hypothesis for the thalamus, Francis Crick has proposed that the relay cells in the thalamus may be involved in visual attention, and has provided an explanation for how this could be accomplished based on the anatomy of the thalamus. This searchlight model of attention and other hypotheses for the function of the thalamus are being explored with computational models and new experimental techniques are being used to test these models. The thalamus is also highly active during sleep. Detailed models of thalamic networks can reproduce the low frequency oscillations observed during sleep states, when feedback connections to the thalamus affect the spatial organization of the rhythms (Steriade et al., 1993) (see animation [1] *– on CD-ROM version*).

Although the spike trains of cortical neurons are highly irregular, information may be conveyed in the timing of the spikes in addition to their average firing rate. This has already been established for a variety of sensory systems in invertebrates and peripheral sensory systems in mammals (Rieke et al., 1997), but whether spike timing carries information in cortical neurons is an open research issue. Brain theorists and experimentalists have engaged in an active dialogue that has enhanced both the interpretation of the data and the theoretical predictions of the models (Abeles, 1991).

4. Technology for brain modeling

New technology is needed to scale up simulations from thousands of neurons to millions of neurons. Parallel computers have become available which permit massively-parallel simulations, but the difficulty of programming these computers has limited their usefulness. A new approach to massively-parallel models has been introduced by Carver Mead, who builds subthreshold cMOS VLSI (Very Large Scale Integrated) circuits with components that directly mimic the analog computational operations in the brain. Several large silicon chips have been

[1] The animation shows spatial patterns of burst discharges in a model of thalamic oscillations. 50 thalamic relay cells and 50 reticular nucleus cells with reciprocal connections and organized in a one-dimensional array were simulated using potential for each neuron was coded using a color scale ranging in 10 steps from -90 mV (blue) to -40 mV (yellow). The ionic channels in these neurons produce bursts of fast sodium action potentials. The activity consisted in a series of distinct clusters of activity propagating in the same direction. In these simulations, bicuculline-induced oscillations were simulated by blocking the fast $GABA_A$ synapses in the network, leaving the slower $GABA_B$ synapses. For more details about these simulations and the experimental observations (see Destexhe et al., 1996).

built which model the visual processing found in retinas. Analog VLSI cochleas have also been built that can analyze sound in real time. These chips use analog voltages and currents to represent the signals, and are extremely efficient in their use of power compared to digital VLSI chips. A new branch of engineering called *neuromorphic engineering* has arisen to exploit this technology.

Recently, Misha Mahowald and Rodney Douglas designed analog VLSI chips that mimic the detailed biophysical properties of neurons, including dendritic processing and synaptic conductances. This has opened the possibility of building a "silicon cortex" (Douglas et al., 1995). Protocols are being designed for long-distance communication between analog VLSI chips that use the equivalent of all-or-none spikes, the same way that long-distance communication between neurons is accomplished. Many of the design issues that govern the evolution of biological systems also arise in these neuromorphic systems, such as the trade-off in cost between short-range connections and expensive long-range communication. Computational models that quantify this trade-off and apply a minimization procedure can predict the overall organization of topographical maps and columnar organization of the cortex.

5. Conclusions

Although brain models are becoming increasingly accepted into neuroscience as tools for interpreting data and generating hypotheses, we are still a long way from having explanatory theories of brain function. For example, despite the relatively stereotyped anatomical structure of the cerebellum, we still do not understand its computational functions. Recent evidence from functional imaging of the cerebellum suggests that the cerebellum is involved in higher cognitive functions and is not just a motor controller. Modeling studies may help in exploring these competing hypotheses. This has already occurred in the oculomotor system, which has a long tradition of using control theory models to guide experimental studies.

At this stage in our understanding of the brain, a model should only be considered a provisional framework for organizing thinking. Many partial models need to be explored at many different levels of investigation, each model focusing on a different scientific question. As computers become faster, and as software tools become more flexible, computational models should proliferate. Michael Arbib and his colleagues have developed an integrated system called Brain Models on the Web (BMW) that allows experimenters to access a variety of models and modelers to access experimental data over the World Wide Web. Close collaborations between modelers and experimentalists can be facilitated by the internet in ways that we are just beginning to appreciate.

Acknowledgements I am indebted to many discussion with colleagues, especially with Christof Koch, Patricia Churchland, and Francis Crick. Geoffrey Goodhill and Venkatesh Murthy made helpful comments on the manuscript. Preparation of this chapter was supported by the Howard Hughes Medical Institute.

References

Abeles M (1991): *Corticonics: Neural Circuits of the Cerebral Cortex.* Cambridge, England: Cambridge University Press
Arbib AM (1995): *The Handbook of Brain Theory and Neural Networks.* Cambridge, MA: MIT Press
Churchland PS, Sejnowski TJ (1992): *The Computational Brain.* Cambridge, MA: MIT Press
Destexhe A, Bal T, McCormick DA, Sejnowski TJ (1996): Ionic mechanisms underlying synchronized oscillations and propagating waves in a model of ferret thalmic slices, *J Neurophysiol* 76:2049–2070
Douglas R, Mahowald M, Mead C (1995): Neuromorphic analogue VLSI. *Annu Rev Neurosci* 18:255–281
Koch C, Segev I (1997): *Methods in Neuronal Modeling: From Synapses to Networks*, 2nd edn. Cambridge, MA: MIT Press
Rieke, F, Warland D, de Ruyter van Steveninck R, Bialek W (1997): *Spikes: Exploring the Neural Code.* Cambridge, MA: MIT Press
Segev I, Rinzel J, Shepherd GM (1995): The theoretical foundation of dendritic function: Selected papers of Wilfrid Rall with commentaries. Cambridge, MA: MIT Press
Steriade M, McCormick DA, Sejnowski TJ (1993): Thalamocortical oscillations in the sleeping and aroused brain. *Science* 262:679–885

See also Integration, neural; Computer and brain; Neurometrics; Hodgkin-Huxley analysis; Information theory and brain function; Artificial intelligence; Neural nets, modeling; Neural networks and neural computing

Relevant website http://www.cnl/salk/edu/CNL

Computer and brain

Michael A. Arbib

The question "Is the brain a computer?" is the modern version of the age-old question "*Is man a machine?*" (see Appendix A) which has played a large role in Western philosophy since Descartes held that the beasts were automata but that human beings had in addition a soul which communicated with the body-automaton via the pineal gland. Today's dualist may argue that the pineal gland has been supplanted by the supplementary motor area, but present-day monism is driven by analogies between brain and computer. The current concern with the problem may be seen as an outgrowth of the rise of *cybernetics* (see Appendix B) in the 1940s, spurred by the development of programmed electronic computers and of sophisticated servomechanisms. In his *Cybernetics, or Control and Communication in the Animal and the Machine*, Norbert Wiener (1948) related the cybernetic metaphor to earlier metaphors, such as the comparison of the body to a steam engine which eventually gave rise to the science of physiology. The study of physiology did not so much reduce the body to extant machines as develop the vocabulary of energy flow and balance in machines to address a wide range of bodily phenomena. A similar transition is occurring with respect to brain and computer.

In the early days of cybernetics, dominant brain-machine analogies included the study of feedback in control of movement, and the logical behavior of automata composed of nets of neuron-like elements (McCulloch and Pitts, 1943). By the late 1950s, however, the interdisciplinary study of animals and machines began to give ground to specialized subfields. Workers in automata theory proved theorems about Turing machines or formal networks with little regard for psychological or neurophysiological validity; workers in artificial intelligence (AI) developed programs to play chess or checkers, to prove theorems or to solve word problems, with few constrained by psychological data and almost none by data on the brain. In recent years, workers in AI have joined forces with linguists and cognitive psychologists to form a new interdisciplinary grouping, *connectionism* (see Appendix F), in which the emphasis shifts from brain–machine comparisons to mind–machine comparisons, seeking to explain human symbolic behavior in terms of mental operations implemented on abstract neural networks whose neurons have no necessary relation to identified biological neurons. By contrast, *computational neuroscience* seeks to benefit from the insights of neural network modeling without losing concern for the data of neuroscience. Here the neurons, or other structures of greater or lesser resolution, are directly related to structures observable in living brains. To understand this approach, we need the notion of "simulation" and the distinction of "top-down" from "bottom-up."

A *simulation* is something which presents the external appearance of something else. In computer science, a simulation is more particularly a computer program which can generate a numerical or symbolic representation of something else. Thus, in *compartmental modeling* (see Appendix D), a computer can simulate a single biological neuron when equipped with a program for solving a suitably parameterized family of generalized Hodgkin–Huxley equations, with one set of equations to describe each compartment (i.e., a region of the cell of roughly uniform electrical activity). But no one would claim that a computer so programmed is a neuron. A computer programmed with Newton's laws can simulate movement but not exhibit it. However, a *robot* (see Appendix E) – a computer equipped, say, with TV cameras for eyes and with mechanical limbs – not only simulates movement, it actually moves. It is still a question of heated debate as to whether an AI system which can take a story and questions about it as input and deliver "intelligent" answers as output is really intelligent or merely *simulates intelligence* (see, for example, the discussion of Searle's "Chinese Room Problem" in Arbib and Hesse (1986)). My own response is an evolutionary metaphor: just as amoeba-like creatures evolved into humans, so current machines may evolve into indisputably intelligent machines; it may then be a terminological choice as to whether a current AI system is a little intelligent or simulates some aspect of intelligence. However, taking a leaf from the Husserlian position of Dreyfus (1979), one should not expect such "intelligence" to be a *human* intelligence if exhibited by a machine that lacks a human body or human social interactions. Having thus sampled the debate on "Can a human-built machine be intelligent" we now return to the question "Is the brain a computer?"

Whether a monist or a dualist, the brain theorist seeks to show that much mental activity or other aspects of animal behavior – such as action, perception, and memory – can be explained by brain mechanisms. Where the connectionist is content with an artificial neural network which simulates (and possibly exhibits) the overt behavior of a human engaged in cognitive activity, a brain theorist seeks a structural as well as functional homology: the program should simulate the interactions of subsystems comparable to anatomically defined regions of the brain – be they the large regions corresponding to the data of the neurological clinic or functional human brain imaging, intermediate regions like layers, columns, or modules, or individual neurons or *subneural components* (see Appendix H). He may proceed *top-down*, seeking to find a plausible functional decomposition of some overall behavior (using, e.g., the tools of *schema theory* (see Appendix C) or directly mapping brain regions to models of the neural circuitry which comprises them), or *bottom-up*, seeking to find the properties of various interconnections of low-level components. In the end, these analyses meet in the interaction of theory and experiment. The brain at these levels of analysis is not like a serial computer executing a single program; rather, the style of the brain is *cooperative computation* in which overall functions are subserved by the concurrent, mutually shaping, excitatory and inhibitory activity of many subsystems. This approach to understanding the brain combines results from a "computational" viewpoint on a particular task with data on behavior, psychophysics, anatomy, and physiology, a pattern of interaction which may change continually as a result of *neural plasticity and other adaptive processes* (see Appendix G). The present article is intended to address philosophical issues in relating computer to brain, not to survey brain theory. However, the author has edited a massive *Handbook* (Arbib, 1995) to which the reader is referred for reviews of many different facets of brain theory, as well as (in Part II) road maps which guide the reader to progress in such diverse areas as the study of biological networks, mammalian brain regions, vision, other sensory systems, plasticity in development and learning, motor pattern generators and neuroethology, and primate motor control. Rather than reducing the brain to a conventional machine, the brain theorist may actually provide a powerful set of concepts to help computer scientists come to terms with the new reality of large networks of cheap *microprocessors* (see Appendix H).

The brain is a computer of an ancient and subtle architecture that transcends the reach of our current technology but inspires new developments in computer architecture. The same *Handbook* (Arbib, 1995) thus also reviews the attempts to build a practical technology for adaptive parallel computation using *artificial* neural networks whose neurons are rather loosely modeled on models of biological neurons but whose interconnections can be modified by a variety of *learning rules* (see Appendix G), most notably those of Hebbian plasticity, perceptrons, backpropagation and reinforcement learning. *Handbook* papers on these themes are grouped under such topics as learning in artificial neural networks, computability and complexity, control theory and robotics, applications of neural networks, and *implementation of neural networks* (see Appendix H).

Appendix A

In 1748, in *L'Homme Machine*, La Mettrie suggested that such automata as the mechanical duck and flute player of Vaucanson indicated the possibility of one day building a mechanical man that could talk (see La Mettrie, 1953). However, although these clockwork automata were capable of surprisingly complex behavior, they were unable to adapt to changing circumstances. In the following century, machines were built that could automatically counter disturbances to restore desired performance, such as Watt's governor for the steam engine, which would let off excess steam if the velocity of the engine became too great. The paper "On Governors" by Maxwell (1868) aid the basis for both the theory of negative feedback and the study of system stability.

Appendix B

Section I.2, "Levels and Styles of Analysis", of Arbib (1995) presents the interdisciplinary nexus in which the study of brain theory and neural networks is situated. The attempt to understand the mind and to build intelligent machines includes, but is in no sense restricted to, the study of neural networks, and so the section begins with a historical fragment which traces our federation of disciplines back to their roots in *Cybernetics*, the study of control and communication in animals and machines. I look at the way in which the research addresses brains, machines, and minds, going back and forth between brain theory, artificial intelligence, and cognitive psychology. I then review the different levels of analysis involved – whether we study brains or intelligent machines – and the use of *schemas* (see Appendix C) to provide functional units that bridge the gap between an overall task and the neural networks which implement it.