The Computational Brain.
1992
Patricia Churchland & Terrence Sejnowski

Chapter 1. Introduction

“The working hypothesis underlying this book is that emergent properties are high-level effects that depend on lower level phenomena in some systematic way.” pg 2

Two groundbreaking discoveries in the 19th century established the foundations for the nervous system 1) macro effects displayed by nervous systems are dependent on individual cells with axons and dendrites and 2) these cells are essentially electrical devices.

“... knowledge of the molecular and cellular levels is essential but, on its own is not enough” pg 4 for example: the stomatogastric ganglion of the spiny lobster contains about 28 neurons and drives the rhythmic action teeth of the gastric mill but it is not understood how the cells interact to constitute the rhythmic pattern!

“as we use the term 'computational neuroscience' aims for biological realism in computational models of neural networks....” p 6

Broad clues .... pg's 7-10 ..... 1) “.... the brain appears to be an interconnected collection of special-purpose systems that are very efficient at performing their tasks but limited in their flexibility” 2) “... the clues about the brains computational principals that can be gleaned from studying its micro-structure and organization are indispensable to figuring out its computational organization because the nervous system is the product of evolution, not engineering design.” 3) “human nervous systems are by no means exclusively cognitive devices.... .... neither the nature of the computational problems the nervous system is solving nor the difficulty of the problems confronting the nervous system can be judged merely by introspection” 4) “of critical importance in generating hypothesis in computational neuroscience concerns the time available for performing the computation .... [and] the amount of space available ..... [and] power consumption ..... [and] the materials available for construction”

Problems in neuroscience have tended to be conceived at various levels which are evolving in parallel, but not always with compatible ideas.

“Data rich, but theory poor” is a description frequently applied to neuroscience.

Chapter 2. Neuroscience Overview.

This chapter encompasses a broad and clear summary of neuroscience as it relates to computation – worth a re-read as an overview.

Levels in the nervous system:

Levels of analysis: Marr (1982) articulated three independent top-down levels of computational analysis (computational, algorithm and implementation) but Marr himself was heavily influenced by neuroscience (a bottom up approach).
Levels of Organization: There are many levels of scale in the brain from systems down to molecules each dependent and intertwined with the other.

Levels of processing: The visual system has roughly hierarchical levels of primary processing (with generous feedback) until it reaches the cortex and fans out. The time for a stimulus to ascend to the cortex and for the motor response to descend again makes up the vast majority of reaction time with relatively little time remaining for the response to be assembled in the cortex. (hints here - pg 25 - of single neuron local knowledge reminiscent of Lamb)

Structure at various levels of organization: (currently identified 7 levels are tentative on further research)

Systems: in the brain are difficult to isolate and define as they are densely interconnected and interdependent. There are almost always reciprocal connections between brain areas. A large number of neurons are almost always involved in the activation of a “downstream” neuron.

Topographic maps: are identifiable for sensory and motor systems for most areas of the body where adjacent areas of the body hold adjacent locations in the brain. (one possible exception is the olfactory system). The brain is structured in part by competitive forces which continue to some extent into adulthood. The cortical surface devoted to a particular body area can shift as much as 1-2cm in the weeks following injury or hyper-stimulation.

Layers and columns: Laminar organization is common in the brain and the cortex is both laminar and functions are arranged in columns thru the layers. However the columns rarely have sharp boundaries and more typically (not always) vary continuously across the surface. Spacial proximity may be a more general rule which may place information required for a particular task in close proximity.

Local networks: each mm^3 has approx 10^5 neurons an 10^9 synapses. These local networks are difficult to study as they are extremely tangled and most knowledge is from recording single units in the network.

Neurons: there are many many different types of neurons (est 50-500) based on anatomical and chemical differences but the distinctions are somewhat arbitrary.

Synapses: chemical synapses are a basic unit of structure which have been highly conserved during evolution but the variety of synapses are highly specialized and knowledge of this are is changing rapidly.

Molecules: membrane proteins which serve as ion channels can be voltage sensitive, chemically activated or both. These in turn can affect propagation of a signal down the length of an axon or neurotransmitter release at the presynaptic terminal. Axon membrane typically contains channels and conductances which permit it to spike when depolarization reaches a certain threshold. "Electrical signaling in neurons is achieved by ionic currents which are regulated by ion channels and ion pumps in the cell membrane. Signaling between neurons is mediated by neurotransmitter receptors in the postsynaptic membrane that respond to particular neurotransmitter molecules by transiently and selectively changing the ionic conductance of the membrane. There are also receptors along the membrane outside the synaptic site that appear to be functional but their role is not known" pg 48

A Short List of Brain Facts: There is specialization of function in different regions of the nervous system. Nervous system has approx 10^{12} neurons and 10^{15} synapses (approx 1 synapse per μm^3). Each cortical neuron is connected to approx 3% of the neurons underlying the surrounding square mm and most connections are between different classes of neurons. A neuron has analogue inputs and discrete a output when a threshold is exceeded. The brain must operate in real time and coordinate different movements in time. An action potential lasts about 1 msec and synaptic transmission about 5 msec. Synaptic potential can last from 1msec to many minutes. Transmission velocity in myelinated axons is about 10-100 m/sec and in unmyelinated axons less than 1 m/sec. The effect of one synaptic input is typically about 1-5% of threshold required. Different neurons have different firing patterns (regular spiking, fast spiking and repetitive bursting) from 3-50 spikes per second. Somatosensory receptive fields vary for different parts of the body and receptive fields in primary areas are for smaller areas and higher areas integrate larger areas. There is a center/surround organization in visual systems where the center area excites and the annulus inhibits and vice-versa. There are also more advanced and specific center/surround effects. Receptive fields are fairly dynamic with repeated stimulation enlarging the receptive field. There are also five sources of non-specific projections each associated with a particular neurotransmitter (brain stem - norepinephrine : raphe nucleus in the midbrain – serotonin : substantia niagra – dopamine : nucleus basalis in the basil forebrain – acetycholine : mamillary region of the hypothalamus – GABA). Some neurotransmitters may be dumped into the intercellular space or the circulatory system for a wider effect. Highly parallel.

Chapter 3: Computational Overview.

This chapter contains a summary of computation developments to date which is good for comparison and appreciation but short on detail.

What is computation? The serial digital computer should be considered a special case and a computer is a physical system with physical states is in some way representational of some other system.

A non-computable function is one where there is an infinite number of pairs for which no rule can be provided (in principle). A non-linear function is one that doesn't map to a line. A vector is an ordered set of numbers. A state space can be represented as a trajectory in state space.

A system computes some function when there is a systematic mapping the sequence of internal states executes an algorithm for the function.

We count something as a computer iff there are functions to describe the behavior of the system and we care about what the function is.

Lookup table: is simple but constrained by memory (exponential for real world problems) and lookup time. "...neural nets may have some properties akin to lookup tables" pg 76 and close but imperfect matches may often suffice.

Linear Associators: have a single input layer of nodes connected to a single output layer. The output is a linear transformation of the input.

Auto-associative content addressable memory (vector completion task) can be achieved by having the same number of input and output nodes and use the Hebb rule (connection weights should strengthen as a function of the correlated activity of the two units) to update the weights. These networks cannot associate two different vectors. The performance of these networks degrades rapidly as the number of stored vectors increases. Modifications to improve performance include making the vectors sparse, normalize the incoming by inhibitory connections, and to introduce a non-linear threshold for output units.

Constraint satisfaction: Hopfield networks and Boltzmann Machines. Figure-ground segmentation (ie necker cube) cannot be performed by a lookup table because it requires the system to have a global perspective while it does not have a global unit. This section considers constraint satisfaction by a process of relaxation via parallel processing.

Linear associators have a feed forward topology while constrain satisfaction require the information to circulate which is achieved using feedback and non-linearity.

Two difficult problems 1. how to adjust the weights and 2. will it find a stable config?

Hopfield proved in 1982 that a particular class of feedback networks will always converge. On hopfield nets, only transitions to states with lower energy are allowed which guarantees convergence to a local minimum. Kirkpatricked invented the technique of simulated annealing for global optimization and these became know as Boltzmann machines in recognition of Ludwig Boltzmanns contribution to statistical mechanics and the networks dynamical properties.

Hopfield and Tank moved from the binary values of a Boltzmann machine (with probabilities of being 0/1) to units with continuous values. The nets do not necessarily find the best solution but they often good solutions very rapidly.

Learning in neural networks my involve weight setting in a supervised or unsupervised context. Supervised learning relies on the input, the nets internal dynamics and an evaluation of the quality of the weight setting. Unsupervised learning involves only weight setting and the dynamics of the network. Supervised learning involved feedback which is external to the network while internal feedback is referred to as monitored.

The problem for hidden units is to discover what combinations of features are ignorable, which systematically occur together and among these which to "care about. Unsupervised learning is good at finding combinations but cannot know what to "care" about. Vise versa for supervised.

If units are linear then there is an optimal solution called principal analysis but many if the "interesting" structures in the world require higher order properties for recognition and analysis. Hence procedures more powerful than principal analysis must be found as trial and error is hopeless in the face of many hidden layers and non-linearity. A procedure is known for Boltzmann machine which involves involves clamping the inputs, wait for equilibrium and compute co-occurrences, then unclamp the inputs and wait for equilibrium, subtract the two and adjust the weights accordingly. Once this approach was found and proved others were sought and found.

Competitive learning: One unsupervised learning approach is to treat the frequency of a sensory input as an indication of value determining a category.

Mutual inhibition can create a “winner take all” setup but it requires normalization to prevent one node increasing without bound. Three problems: sometimes critical information is not the most critical, this approach picks out only the lower order features, and the weight setting tends to be unstable with novel input. (it is not understood how nervous systems cope with this problem)

Curve fitting: is a classic example of fitting parameters to a model. When the parameters get large the algebra becomes nasty. A gradient decent approach is one solution but it has local minma problems in the non-linear situation.

Feed-forward nets: have no feedback but a speed advantage. They cannot do XOR. The trick is to put the output of the hidden units through a squashing function – a smoothly varying nonlinearity such as $\sigma(x) = 1/(1+e^{-x})$. A feed-forward net with an sufficiently large number of hidden units can approximate any well behaved function. However, they do not scale well.

Discriminating sonar echoes: of a rock vs a mine can be done as well as a human with a feed forward network.

NETtalk: was a system successfully trained to produce English words from phonemes.

Recurrent Nets: Feedback allows the network to incorporate multiple time scales, to process temporarily extended input sequences, to generate oscillations and rhythms of varying freq, to resolve ambiguities such as figure-ground. Recurrent networks can have output iff there is ext input and internal input, have output with an external trigger followed but continuous internal activity, or have output due to internal input alone with external input optional.

Networks may recognize and handle the temporal nature of much real world data. On crude approach is to map a temporal sequence onto a spacial sequence. Or vise versa with a spacial sequence (a word) moved past a “window”. Another related approach is to use feedback to “let time represent itself” in the sense that the network takes time to settle and the continuous input simply interacts with the settling process (appropriate for very short time frames). This is analogous to adding a slow decay into the network to prevent the influence of input being abrupt. These Jordan nets achieve this temporal capability but cannot match a variety of inputs to a variety of outputs.

How can long term memory be added to Jordan networks? Change the time constant, change the architecture, or change the activation function. The choice must be matched to the problem at hand.

A nervous system may not conform to a single blueprint but will likely hit upon different solutions to different problems.

From Real World to toy World: Real world generally has more dimensions (causing scaling problems) and real world inputs do not arrive as identified types. Fattening up a network to cope tends to suffer scaling problems of increasing weights to set and data required to train. A more serious problem concerns how information is distributed in the network in that new training data disrupts all weights and even worse, when a weight changes the gradients require recalculation.

Radial basis functions allow hidden units to be sensitive to only a limited range of new data. An advantage is that training only impacts part of the network. Radial basis functions require that the centers and widths must be chosen. In low dimensional data one trick is to uniformly cover the space with spheres of a given size (in high dimensional spaces the # spheres required can be very large). A related solution decouples the network into smaller subnets. A more sophisticated solution has the net determine how to cluster the subnetworks. And a more sophisticated approach again will apply competition between subnetworks (Nowlan 1990, Jacobs et al. 1991). These approaches utilize a referee net which directs the data to the correct subnet. This achieves some measure of success and has some biological parallels in bilingual people.

What good are optimization procedures to neuroscience? “Assuming the architecture is relevantly similar to the anatomy, and assuming the dynamics are relevantly similar to the physiology, then so long as both the actual net and the model net use some parameter-adjusting procedure to minimize error, then they will each end up at locations on the error surface that are rather close. How close the final positions are depends on how similar the architecture and dynamics are, but the critical point is that error minimization is an optimization procedure, so it is reasonable to expect that however a given net finds its error minimum, it will be as close to the error minimum found by another net as long as the architecture and dynamics of the two nets are closely similar.” pg 131

Is the brain in the error minimization business? 1)There are too many synapses to be set by DNA 2) During development there is massive competitive cell death suggestive of optimization 3) The nervous system uses feedback 4) Natural selection culls outperformed nervous systems.

Models realistic and abstract: No model is 100% realistic. What goes into the model depends on what one is trying to explain, and in the nervous system, that is intimately related to the level of organization one is targeting. Make the model simple enough to reveal what is important, but rich enough to include whatever is relevant to the measurements needed.

Concluding remarks: Lookup tables may well be used in some functions of nervous systems where there response time is at a premium and training may sometimes be the process of moving a novel situation into a lookup table.

Chapter 4. Representing the World.

Distinguish between a current representation (such as perceiving or imagining) and an abeyant representation that is part of one's background knowledge.

Constructing the Visual World: The human visual system has some 10^8 transducers but only approx 10^6 axons leave the retina for the brain. About 10^{10} neurons are involved in visual processing beyond the periphery. Signal integration and processing occur at every level to somehow build full-blown perceptions but it is not built “brick by sensory brick”.

Genes form the general contours of the brains capacity. Human visual systems generally assume that objects are lit from above. Ramachandrans hypothesis states that the visual cortex works according to assorted rules of thumb – inchoately specified genetically, stabilized during development, defeasible computationally and sometimes by top-down influences.

It is important to understand both the animals behavioral repertoire and the specific range of physical parameters that a particular species uses to inform and initiate behavior. What an animal represents a perception depends on the animals environmental niche and how it contributes to survival.

Thumbnail sketch of the mammalian visual system. This is indeed a thumbnail sketch and difficult to further summarize – well worth re-reading pg 148-157. Search terms retina, rods, cones, ganglion, color, motion, contours, location, shape, direction, speed, reflectance, depth, lateral geniculate, thalamus, parvo, magno, V1, V2, V3, superior colliculus, ocularity,center-surround.

Representing in the brain: what can we learn from the visual system? Partly by accident, Hubel and Wiesel discovered (1962) response selectivity of neurons in visual cortex to a light-dark edge in a given orientation – close cells respond to slightly different orientation. These results prompted a lot of single neuron research but there was no comparable breakthrough regarding cell assemblies. A theory of information processing began to condense around 4 points 1) cellular response selectivity 2) a processing hierarchy 3) fewer cell responding at higher points in the hierarchy 4) segregation of parallel pathways for different visual stimulus (i.e. color, & motion). Local coding and grandmother cells was promising but ultimately failed. The (preferred) vector coding alternative has superseded local coding as it 1) is just as consistent with data 2) is better at recognizing and responding to new things and 3) recognizes that the number of patterns recognized outstrip the number of sensory processing neurons.

What is so special about distribution? The differences between local and distributed representation may appear trivial - the grandmother cell is not really realistic, local vs distributed are not mutually exclusive – but the distinction is important…. **What is vector coding?** Parallel processing is possible without distributed representations. Distributed representation can be course or fine coded. Redundancy improved signal to noise ratio, amplifies a low-intensity signal and provides backup. **The conceptual fecundity of state space**. Weight space dictates the partitioning of activation space. Weight configurations both cluster different things and may be sensitive to very fine differences. An activation space is a similarity space in that nearby objects are similar – an intrinsic feature. Network concepts encapsulate both representation and processing. The vector coding/vector-matrix processing framework makes roughly comprehensible the ease which humans extend concepts to a new member of a class. Network computation is robust in the face of damage to the network or to input data. (local coding has no equivalent framework for these instructive aspects of the distributed model)

World enough and time. We are discussing simplified networks and the questions of dynamics has been largely absent. There are a variety of processes in real neurons that are neither clearly activations or weight modifications. "Evolution it is evident, exploits whatever odds and ends are handy in order to fit out the system to thrive…" pg 176. "Speed of response is now the only thing that distinguishes activation from weight modification, and it does so relatively, not absolutely". Pg 177

What is course coding? Course coding can achieve hyper-acuity with low resolution transducers.

Does tuning selectivity of neurons imply local coding? "Hierarchical organization, pathway segregation and tuning selectivity in neurons are consistent with, but do not entail, local coding. They are also consistent with vector coding." Pg 179 For a node to have zero activation is not to have no value – it is to have the value 0 – which means "something". The fact that a neuron responds or does not respond to a stimulus is not a sound method of determining cause and effect. Local coding is limited to the number of nodes while vector coding has very much higher limits.

Shape from shading: A neurocomputational study. One of the primary properties of a surface is its curvature. Directions of minimum and maximum curvature are always at right angles to each other and are called principal curvatures (and provide a complete description of the local curvature). Principal curvatures are difficult to extract from gray-level shading due to many factors, but our visual systems do. A neural model shows that the hidden units have receptive field properties similar to those of similar cells in the visual cortex (edge detection and bar-detection). Further, there is no way of determining the function of each hidden unit simply by examining the receptive-field properties of the unit.

Stereo vision. **Psychological Parameters**. Humans are usually unaware of stereoscopic vision, motion parallax and color constancy. Good 3D perceptions diminishes beyond 100m. Images from each eye are aligned at the vergence zone and objects within 10-20 min of arc fore and aft of the vergence zone are slightly disparate but merged by the brain (Panum's fusion area). Eyes make vergence shifts as often as every 200ms. Stereo perception is independent of object identification (eg random dot stereograms). About 10% of humans are stereo deficient – but are rarely aware of the deficiency. **Neurophysiology and Anatomy**. Striate cortical cells have ocular columns which are bands of cells which are strongly monocular, weakly monocular and binocular. Very few of the synapses in V1 and V2 originate in the LGN, rather the connections are mainly inter-cortical. In the striate cortex there are tuned cells (respond to zero disparity), near cells (respond to objects fore of vergence), far cells (aft of vergence) and a forth more general class of cell. Relative depth cells typically also respond to movement in a particular direction and orientation. Cells responding further from the vergence tend to have wider tuning curves.

Computational Models of Stereo Vision. Matching the image on one retina with the other (the correspondence problem) is difficult because the corresponding points have discontinuities, there may be elements on the left which do not exist on the right, uniform expanses are problematic, correspondence info is sparse requiring interpolation and there will be noise. If course stereo is the first processing step then how far can a feed-forward net go, is it biologically plausible, how can noise be cleaned and when should other depth cues be introduced? **Finding compatibilities and course stereo**. The fusion-net model uses an array of Not-XOR nodes to locate compatibilities on the vergence depth and two additional arrays for fore and aft. This configuration indicates the match at the vergence depth but also provides information about for and aft matched by means of units which are displaced just one position to the left or right. The relations between the three output layers provides information about which match is fore and aft as well as size, shape and position cues. The network is fast (feed-forward) to run, and learn (2000 trials), provides data from low res to guide subsequent high res, mimics high resolution in Panum's area, and the matched move with the fixation depth. The architecture will scale and is biologically plausible. Mammalian systems perform feature analysis after the pathways from both eyes have been converged. The technique make evolutionary sense for camouflage breaking. Also infants develop hyper-acuity at about 8 weeks with control over eye movements in order to verge and fixate (and precedes other depth cues). Noise suppression might be effected with mutually inhibitory connections between the three classes of output cells. **Using a Cooperative Algorithm**. Marr and Poggio developed a basic strategy of combining a matching function with a cooperative procedure with a 3D grid representing multiple depths with matches at the same depth interacting positively and at different depths interacting negatively. However it is biologically implausible because humans do not experience crisp images at all depths – the mechanism fails the "screen door" test. **Unsupervised Learning of Stereo Invariants**. This is difficult for a feed forward net because representing disparity as the invariant is a second order property. However, the problem can be solved by using anti-Hebbian connections (ie the weights to the winner are decreased) and replacing the sigmoid with a Gaussian function (large output from a balance of inputs rather than maximal input). **Depth cues and the Real World**. Random dot stereograms are highly artificial. Most cortical neurons in visual cortex respond best to edges, borders, motion and disparities. This preprocessing essentially accentuates any discernible features. In area MT single neurons are tuned for both disparity and relative motion and in area MST (receiving input from MT) have direction selectives that depend on the distance of the moving object from the plane of fixation. Single processing are probably not taken to completion by integrated as a whole.

Hyper-acuity: from Mystery to Mechanism. Involves detection of intervals smaller than the that possible with a single transducer. **From three cones to 10,000 hues by Course Coding**. Red looks red under a wide variety of lighting conditions (color constancy). Color perception depends on wavelength but is not identical to wavelength. Rods have peak sensitivity at 510nm (but are not involved in color perception. Cones have broad response curves with peak sensitivity at 420, 530 and 560nm: known as B,R & G respectively but actually look more like violet, blue-green and yellow-green. Light usually has a broad power curve rather than a single wavelength. Any perceivable difference corresponds to a distinct vector tripple. **Vernier style hyper-acuity**. … **Hyper-acuity in Depth Perception**. Humans have depth discrimination with an accuracy of about 5 arc sec which falls of rapidly away from the plane of fixation (the tuned cells at the point of fixation are very narrow, getting broader with further "distance" from the zone of fixation). Accurate depth perception requires more than the three cell types used in vision. Investigation of a model showed that the tuning curves further from the plane of fixation contributed more to the depth discrimination (as it is the slope of the tuning curve that is important). **Vergence and absolute depth judgments.** A model demonstrated that eye angle and disparity can be combined to generate an egocentric depth (a reminder that the same cells can be used for different purposes). In addition, the hidden units developed a response to disparity gated by vergence (ie they had acquired "gain fields" meaning that the units response was dependent on eyeball angle.) Neurons have been identified with gain fields

Vector Averaging. The nervous system may use different techniques in different areas. Vector averaging may be used in the superior colliculus and the motor cortex. Color and depth discrimination involve multi-dimensional representation. In contrast, vector averaging reduces dimensionality and makes sense for the likes of motor output where a single outcome is required. An example is eyeball position changes which are relative to the current position.

Chapter 5. Plasticity: Cells, Circuits, Brains and Behavior.

Introduction. Neuroscientific research on plasticity roughly divides into four main streams: 1. the neural mechanism for simple plasticity such as classical conditioning or habituation 2. studies of temporal lobe structures including hippocampus, perirhinal structures, and amygdala 3. study of development of the visual system 4. the relation between the animals genes and the development of its nervous system.

This book focuses primarily on hippocampus etc.

Learning and the Hippocampus. Patient HM prompted the focus on the hippocampus and demonstrated the major division in the nervous system between memory for individuals and events which is reportable and other kinds of plasticity such as conditioning, skill acquisition, and recognition without explicit awareness. Patient Boswell later suffered damage causing a 40sec memory. Frantic research of the hippocampus has followed at all levels. “A unique property of the hippocampus is its strategic location as a convergence region for information from nearly all high-order cortical areas, as well as from the brain stem nuclei. Every sensory modality projects to the hippocampus (via entorhinal cortex), and most have reciprocal connects back.” Pg 248

Donald Hebb and Synaptic Plasticity. “The crux of his insight, slightly reconstructed is this: co-activation of connected cells should result in a modification of weights so that the probability of the postsynaptic cell firing given the presynaptic cell fires, is increased.” Pg 251. The simplest form considers pre & post synaptic firing rates (only allows strengthening), one modification is to use the deviation from the average firing rates (allowing positive and negative weight changes) while a third might consider the rate of change……. In Hebian plasticity, the plasticity is 1. specific to the synapse in question, 2. depends conjointly on both presynaptic and postsynaptic cells and 3. depends exclusively on those cells and not others. Various Non-Hebbian plasticity has been identified: activation of A-B strengthens C-B (hetro-synaptic potentiation), activity in A, while inhibited by C, influences A-B, and where potentiation between A-B also influences A-X.

Memories are made of this: Mechanisms of Neuronal Plasticity. Long term Potentiation LTP was discovered in 1973. Anatomy of the CA3 field is reminiscent of a recurrent net with each Granule cell projecting mossy fibers to only about 14 pyramidal cells in the CA3 field; and with each pyramidal cell in the CA3 field innervated by about 46 granule cells. These connections are primarily to the upper reaches of the dendritic tree of the pyramidal cells while the lower tree is dominated by connections from other pyramidal cells in the CA3 field.

[insert fig 5.11 here].

Evidence was found to link LTP to behavioral learning and the molecular basis was established involving Glutamate, glycine, norepinephrine and critically the NMDA receptors. **LTP and Cell Populations**. Population wave forms are the result of synchronization of synaptic potentials (not action potentials) with theta (4-8Hz low amplitude) occurring during exploratory activity and REM sleep, and sharp waves (irregular 0.02-3 Hz high amplitude) occurring during relaxation and deep sleep. Buzsaki (1989) proposed a two step routine “(1) theta rhythm: explore and experience (online); (2) sharp waves: rest and consolidate (off line).” Pg 267 **Beyond vanilla LTP**. Dendrites have processing capacity which is not classically recognized but massively expands brain capacity. Non-Hebian,LTP can be initiated by 5Hz oscillating current (ie theta) dependent on NMDA receptors (and possibly attention related). LTP at the molecular level. Ca^{2+} is a central player and is buffered by calmodulin. There is Ca^{2+} concentration in dendritic spines due to buffering overload and diffusion effects. LTP may be presynaptic, postsynaptic or both. **Extrasynaptic and non-NMDA LTP**. In addition, the threshold is lowered for producing a spike when the signal enters at the potentiated synapse. Another effect in pyramidal cells in the CA3 region appeared to show LTP dependent on non-NMDA receptors.

Cells and Circuits. So what? The hippocampus appears to lie on the path from working memory to long term memory.

Insert fig 5.28 here.

However, the specifics of its function are opaque due to the matrix like structure and the massive back-projection into the entorhinal cortex and from there back into many areas of the brain.

Insert fig 5.29 here

Decreasing Synaptic Strength. Two broad classes of LTD have been postulated; hetro-synaptic LTD (responsiveness of the whole cell is reduced) and homo-synaptic LTD (synapse specific). Note that LTD delivers a theoretical increase to the amount of info stored and does not represent forgetting but tuning. There is evidence to suggest that LTD occurs in a particular synaptic activity “window” just below that required for LTP (ie it is a continuation but with a break-over point).

Back to Systems and Behavior. Damage to the hippocampus does not affect existing long term memories (LTM) or working memory (minutes) but typically destroy the memories in transition/consolidation (last few weeks) from working to long term memory. Interestingly, the "recency effect", see patients with normal LTM but degraded Short Term Memory (STM). Proposed models of STM include a"circulating loop" model and a "transient weight modification" model.

Being and Timing. Time must have pride of place in computational models. Memory (of various lengths) is a mechanism to manage and exploit temporal phenomenon in the environment. (see Traub's model).

Development of Nervous Systems. There is a stunning level of plasticity during development and DNA cannot code for it all. In some structures there is 75% cell death. Case study: development of ocular dominance patterns or columns – wobbly bands in layer 4 of the visual cortex sharing a response-preference for input from either the left or right eye. During early development these bands segregate from random distribution; a process at least partially dependent on input activity from the eye. This process involves local competition of asynchronous stimulus between neurons. Computer modeling has been effective in establishing that "the key factors for the emergence of ocular dominance organization are connectivity, activity correlations between eye represented, and interactions that depend only on the distance between neurons.

Modules and Networks. Computation, representation and plasticity are intimately connected: representations are activation vectors, learning involves changing synaptic weights, and computation is a matter of vector-vector transformations through weights and squashing functions. Nervous systems are vast networks, with various regions specialized for various tasks but with massive recurrence and interdependency. Memory storage is performed by the same structures that perform information processing. Modules are apparent at various levels (cellular, columns, neuroanatomy and neurophysiology) but are difficult to define in a clear sense. Lesion studies are challenged by the variability of lesion size and location, variability of brains and the severely limited access to human brains for research. However, humans can be guided by verbal instructions and can provide verbal reports which offer a rich source of introspective information. An individual representation of a dog "Daisy" is a construct from many sub systems visual, aural etc and sub-sub systems, face recognition, gait, etc and information in any of these subsystem seems capable of activation the whole memory. There is some evidence for inherent categories such as animate vs. non animate, and sub categories within these. "Crudely, there are networks interacting with networks interacting with networks. To a first approximation, recognition defects may be described in terms of a network organization that involves both information convergence and feedback, and relative degree of categorical specificity (A. Damassio)" pg 321. There is a distinction between auditory, visual and olfactory memory, and a further distinction between the sensory recognition capability and the motor action capability. All of these capabilities appear to be intermeshed in a way that is neither clearly modular or hierarchical. "..none of these arguments rule out the possibility of modularity; they imply merely that modularity is not the only, and may not be the best, hypothesis to square with deficit selectivity." Pg 325. There is some evidence of both cross-talk across sub-modalities (i.e. motion versus color) and cross-talk between modalities, or via feedback routes, between higher and lower processing levels.

6. Sensorimotor Integration.

Introduction. The fundamental output of the brain is a motor action. The seeming simplicity of the action would appear to be a poor indicator of the computational complexity involved. For example, the transformation between two eyes and a single arm positioned in different locations, with different degrees of freedom is a complex mapping even before taking into account friction, gravity, inertia, power supply, development, feedback, and time. Study of the simple nervous system of a leech is instructive. The leech has 21 segments, 10^4 neurons, a head and tail ganglia, and a few basic movements such as crawl and bending reflex.

Leechnet. The bending reflex is a very simple escape behavior where the leech withdraws from an irritating stimulus be creating a kink in its body roughly proportional to the size of the stimulus. Invertebrate findings may not map to vertebrates. "...it is astonishing that as tiny a nervous system as that of the leech can be as sophisticated, flexible and successful as it is." pg 341. The reflex is not a direct connection between sense and motor function. **Computer Models of Network for Dorsal Bending**. LeechNet I simulated a single segment with a feed-forward network with 4 input, 8 output and 18 inter-neuron units. It was trained by a supervised learning algorithm, to match the Leach behavior and measured neuron potentials. The weights on individual interneurons are not directly comparable so a statistical comparison was completed which displayed a good match. Leechnet II added some missing dynamic components of neurons and various motor and inhibitory connection not represented in the feed-forward network. The network was trained good match was achieved with the temporal aspects of the Leech. Back-propagation was used as a tool to set the parameters in the network which could then be studied. The studies illuminated some paradoxical neuron measurements (ie confusion between measuring the change in a single neuron and assigning causality vs measuring distributed change) and artifacts that may have been related to integration of this "component" with the rest of the nervous system. Classical conditioning was explored and a larger number of very small distributed changes to units was identified (thought: is habituation a neuron or network feature???)

Computation and the Vestibulo-Ocular Reflex. This section provides a far more satisfying application of network modeling than I have read in earlier texts – well worth a re-read. "Table 6.1 Summary of primate gaze control systems. Hold gaze – fixed target: *Vestibular-ocular reflex (VOR).* A system that uses knowledge of accommodation and vergence state together with head accelerations to stabilize the gaze vector. Hold gaze – moving target: *Vergence.* A binocular system for locking both foveas over the same three-dimensional target. *Pursuit.* A System for tracking moving objects by generating smooth-velocity control signals. Change gaze – *Saccades*. High speed, precomputed movements that rapidly change gaze over small to very large visual angles." Pg 356. VOR is one of the best understood parts of the mammalian brain, with very fast onset (about 14msec) and can recalibrate itself when required (ie glasses). The shortest pathway between transducers and muscles is just a three link chain and normally operates as a feed-forward network.

Insert fig 6.20 here.

"the vestibular nuclei …. is an area of convergence …. Code[d] in a distributed rather than local fashion, … and that [serves] in several different roles, performing either VOR computation, or saccades or pursuit computation." Pg 358. Anastasio and Robinson constructed model networks with inputs of head velocity, saccades velocity and pursuit velocity, and output of eye velocity. Back-propagation was used to train the network and retinal slip was used as an external teacher. The analysis revealed hidden units with a range of pure and combined responses, no single "solution" for hidden units and a tenancy towards excitatory and inhibitory roles (as for real networks). It was notable that as few as two units at each hidden level could solve the VOR problem and that additional units simply distributed the load (redundancy). "Back-propagation enables the network to find an error minimum, and it is not unreasonable to assume that evolution, over many trials and with countless errors, stumbled onto much the same error minimum." Pg 362 The model successfully replicated the networks response to damage (one severed vestibular organ) and provided insight to the nature of the recovery symptoms. **Plasticity in the VOR**. An open-loop system must be recalibrated when the physical properties of the system change (ie during development or eye-glasses). VOR (open-loop, feed-forward) and pursuit (negative feed-back) share an architecture. Retinal slip tells the system to recalibrate. Modeling successfully identified where the recalibration is made by matching the Static and dynamic characteristics of the VOR system. This clarified some experimental results which were at odds with theory, which ultimately assisted in developing a new hypothesis regarding how the recalibration is done. All three systems converge on the one set of eyeball muscles and the VOR has a 14msec delay, the pursuit a 100msec latency and the saccades a ?msec latency. This modeling was interesting because back-propogagation located an error minimum amongst substantial anatomical and physiological constraint which then established that recalibration of the VOR system (due to magnification goggles) mirrored the biological response and was achieved by plasticity in only two specific locations (at the exclusion of seven others). Furthermore, when these two locations were prevented from altering the model did not recalibrate successfully. This makes it highly likely that the model is a stable representation of the real system. Back-propagation is a modeling tool that can be used in place of evolutionary trial and error to establish the global weight profile of a network model constrained by characteristics of the real network. Also " if, in trying to explain a phenomenon such as plasticity in the VOR, one looks at a single location and finds a change, one may be inclined to interpret that as the change that correlates with *the* learning seen in the behavior. But the interpretation may well be misleadingly narrow." Pg 377. In comparing Leechnet and the VOR models; the hidden unit representations, are distributed, carry multiple kinds of signals, and sets of neurons act in concert to produce vector transformations. [see]

Time and Time again. Many organisms achieve a goal in a computationally efficient manner by repeating a simple motor action – but these must be coordinated in time. While many rhythmic actions have an apparently "clockwork" base they are also very often modifiable to circumstance. **Oscillating Circuits in the Spinal Cord**. The spinal cord is an exceedingly complex piece of neural tissue in its own right and is capable of sustaining basic locomotive action in isolation from the brain. However, afferent input is not at all inconsequential to the rhythmic behavior in locomotion as it introduces many timing subtleties and intensity variations to the basic rhythm. The basic model is the Brown-Lundberg half-center circuit.

The Segmental Swimming Oscillator. This section models the swimming action of the Lamprey using the half-center circuit. The action involves a wave of muscle action which travels the length of the lamprey. An important architectural aspect of the lamprey spinal cord is that the bursting output of one segment bears a constant phase lag to the bursting output of the segment in front (about 1% of cycle time), not a fixed delay. The lamprey was successfully modeled with a neural net and mathematical techniques (a chain of oscillators)

Modeling the Neuron. A basic and early summary of the material which is substantially elaborated in the later text by Kosh. It covers Hodgkin-Huxley, Calcium and Potassium currents, channels, dendrite trees etc.

Concluding Remarks. It is important to note that all of these simple reflexes from the Leech to VOR can be modified by signals from higher cognitive level processes.

7. Concluding and Beyond

Almost everything remains to be done! But we need to build synthetic brains operating in the real world to really understand how real brains work.

Mean and Analog VLSI. "Mead's approach is to make analogue integrated circuits with transducers that convert real world signals into electrical representations, and with a design that mimics closely the real anatomy and physiology of the nervous tissue under consideration." Pg 416. Mead and his colleagues have built a synthetic cochlea-brain stem nucleus based on the nervous system of the barn owl.

Ballard and Active Perception. Ballard has shunned the conventional assumption of extensive analysis and categorization and combined three ideas motivated by ethology and evolutionary biology: 1. an animal need to characterize and categorize everything in its perceptual domain, but only what is relevant to its tasks and survival (such perceptions are known as indexical representations), 2. task-relevant perception involves learning, minimally in the form of reinforcement of the perception-action pair when the action provided by the indexical representation is successful, 3. implementing the first two ideas leads to systems that actively control their environment.

Brooks and Mabots. Brooks "inverted the time worn rule in robotics that says, "Simplify the problem and make the robot sufficiently intelligent to succeed flawlessly in the toy-world." [and] … decided the simplify the robot, and see how it could manage to cope with problems in their unvarnished and undiluted state." Pg 423.

Appendix – Anatomical and Physiological Techniques.

An overview is provided of permanent lesions, reversible lesions and microlesions, imaging techniques (CT, MRI, 2-DG, PET) Gross electrical and magnetic recording (EEG, ERP,) and single unit recording (microelectrodes) and Anatomical tract tracing.

Insert fig A.1 here

Further Reading
A number of texts relating to the single neuron are sited as point 6 on page 445.
Nowlan 1990,
Jacobs et al. 1991
Ramachandran 1988, 1990a, b