(11) Grillner, S., Wallén, P.: Central pattern generators for locomotion with special reference to vertebrates.
Ann. Rev. Neurosci. 8, (1985), 233-261.

(12) Selverston, A. I. (ed.): Model neural networks and behavior.
New York: Plenum Press, (1985).

Review:

The Visual System

by *Martin I. Sereno*

Organization of Neural Networks
Structures and Models

Edited by
W. von Seelen, G. Shaw, U. M. Leinhos

Weinheim: VCH Verlagsgesellschaft
1988, pp. 167-184

SUMMARY

I discuss selected aspects of the anatomy and physiology of the visual system that may be of interest to theorists attempting to model real neural circuitry. Though a great deal is already known about the visual system, our knowledge is still quite incomplete, particularly with regard to the principles governing interactions among different types of neurons. I therefore consider some strategies for building stronger links between models and the kinds of data we can now obtain.

1 INTRODUCTION

Neural modelers often construct simple abstract networks out of "units", and then analyze the dynamical behavior of these networks in detail, considering only afterwards their possible relevance to real neural circuits. A similar strategy has at times proved very successful in the physical sciences. A turning point in the development of mechanics at the time of Galileo and Newton, for instance, came with the idea of constructing and studying the physics of a thoroughly abstract and simple "absolute space" lacking friction, gravity, magnetism, and so on. Only after this could the physics of more complex and cluttered real worlds be articulated. Similarly, a neural modeler tries to find the basic mathematical forms (presumably, rather more complex ones) that get fleshed out but also made less obvious in the particular networks that happen to have been patched together in the evolution of real brains. I am in sympathy with this approach and only want to make sure that the basic abstract forms are near enough in complexity to real networks that it becomes possible to begin relating the two. It is often suggested that neural-like "units" need not always be construed as neurons, but instead sometimes as modules or even groups of modules, especially when higher level functions are considered. The use of similar models across several levels of organization, however, would seem to imply the existence of strong between-level isomorphisms unlike anything seen in other biological domains. It is difficult to think, for instance, of any isomorphisms across protein structure, cytoplasmic organization, and multi-cellular assemblies. To avoid glossing over significant between-level differences, I think it is essential to be as explicit as possible in describing how a model relates to biological neurons and neural assemblies. To do this, it may be necessary to start out with more realistic neuron-

like elements and connection schemes than is customary in network modeling. There are several longer reviews covering some of the topics discussed here (1, 2, 3, 4, 5).

2 VISUAL SYSTEM ANATOMY

The visual system, broadly defined, is spread across 40 or more nuclei and laminated structures in the pons, midbrain, thalamus, basal forebrain and cortex. The first four brain regions are often artificially set to the side or completely omitted in discussions of the cortex. This is unfortunate since every visual cortical area has direct access to the motor system via these structures. It seems likely that the nature of the activity in cortical neural networks has a lot to do with this thoroughly distributed sensorimotor linkage. Having said this, I shall nevertheless have to focus on primate cortex for the sake of brevity. My anatomical notes begin with areal features and then consider successively smaller subdivisions.

2.1 Topography of Visual Areas

Primate visual cortex consists of at least 20 distinct, adjoining regions or "areas" that occupy as a group about half of the total area of the neocortex (2, 3, 4, 5, 6, 7, 8, 9). A working list of visual areas in the macaque monkey followed by their probable and possible counterparts in the owl monkey is as follows: V1 (=V-I), V2 (=V-II), V3 (=DM?), V4/VA (=V4?/VA), V4t/VOT (=DL proper), MT (=MT), VP (=VP), PO (=M), V3A (=DI?), MST (=STv), FST (=?), LIP (=FS?), VIP (=PP?), 7a (=TP?), DPL (=?), PITd (=ITcd?), PITv (=ITcv?), CIT (=?), AIT (=ITr?), TF (=ITml?), TH (=ITmm?), TG (=ITp?). This list is incomplete; but the total number of cortical visual areas is certainly no more than 30. The remaining cortex consists mostly of somatosensory, auditory, and motor areas. Cortical areas have extensive direct interconnections within, but not between, modalities. Between-modality connections in monkeys are made in large part via nuclei in the amygdala and areas in frontal cortex (10, 11).

Visual stimuli are tranduced by photoreceptors, processed by 40 to 60 different cell types in each local patch of the retina (12), and passed along about one million retinal ganglion cell axons into the CNS where precise topological maps are established in the superior colliculus, the dorsal lateral geniculate nucleus (dLGN), and in several other smaller brainstem structures. The projection of the dLGN to striate cortex (V1) is as orderly as the retinogeniculate pathway, and V1 has the largest, finest-grained retinotopic map of any visual area; it occupies about 8% of the total neocortex in macaque monkeys (about 17% in owl monkeys, and 1-2% in humans). In primates, the other, so-called "extrastriate" visual areas receive much of their visual input from V1, either directly, or via intermediate cortical areas and also via cortico$_1$-thalamo-cortical$_2$ links involving six or more topographic visual nuclei in the pulvinar (13, 14). (In some mammals [e.g., cats] every extrastriate area receives a direct input from the dLGN). A second source of input is the superior colliculus, which projects to some of the extrastriate-projecting nuclei in the pulvinar and is capable of visually driving extrastriate cortex in the absence of V1 (15). Cortico-cortical connections and lesion studies have defined somewhat separate parietal and inferotemporal processing streams (see next section).

As one moves by stages into the system, retinal receptive fields grow larger and the retinotopy becomes less clear. However, the absolute topological fidelity of different interareal projections does not vary greatly, and much of the increase in receptive field size can be

attributed to the smaller size of most extrastriate areas (from 30% to less than 5% as big as V1) and to the effect of successive slightly divergent remappings. Non-topological projections between adjacent stages in the cortical areas hierarchy appear to be uncommon but have been clearly demonstrated in several feedback projections from areas (CITd, CITv, TF, TH) in inferotemporal and parahippocampal cortex (9). Notably, connections between the thalamic and cortical areas most remote from V1 are topographic (16).

Laminar connection patterns of corticocortical projections can be used to group the 20 or more visual areas into five or six rough hierarchical levels by counting the number of "feedforward" interareal projections one must pass through to get to an area. It is possible to subdivide these levels further. For example, two areas B and C that receive feedforward input from area A might be placed on separate sublevels if B has a feedforward projection to C but not vice versa. In general, there are a number of different feedforward routes to a given cortical area, and these routes often differ in length (i.e., in the number of cortico-cortical projection stages). There are many different routes to higher areas. "Feedforward" projectionssynapse mainly on layer 4 of the target area (the main input layer of cortex--see below), while "feedback" connections avoid layer 4, synapsing above and below it. Most interareal connections are reciprocal. These hundred or so cortico-cortical bundles, each containing a large number of axons, interpenetrate in the white matter below visual cortex, which also contains a hundred or so input and output bundles involving subcortical centers. Limited studies of the latency of response to a flashed stimulus confirm the anatomical hierarchy. Striate cortex is activated with a sharp rise time about 30 milliseconds after the onset of a flashed bar (most of the initial delay occurs in the retina) and each successive interareal feedforward link consumes approximately an additional 10 milliseconds (17, 18) (most of the interareal delay reflects within-area processing). Thus, over 100 milliseconds typically passes before the areas at the top hierarchical levels become active (19).

There are numerous outputs from the visual cortex. The higher areas project into the limbic system. Areas in the top three hierarchical levels project to several nuclei in the amygdala. The areas at perhaps the highest level (TF, TH) project to the entorhinal cortex, which projects into the first of a sequence of areas in the hippocampus (in order: dentate gyrus, CA3, CA1, subiculum). Some of these last projections (e.g., dentate to entorhinal, and CA3 to CA1) appear to be topographic (20). In contrast with the upper level projections into the limbic system, outputs to the motor system emerge from all levels in the hierarchy. Each visual area projects to several of the following: frontal eye fields, caudate nucleus and putamen, intralaminar nuclei (which project to motor cortex), superior colliculus, pontine nuclei. In many cases (e.g., cortical areas to the caudate nucleus and superior colliculus) these projections are topographic.

Some areas contain an almost complete, topologically continuous map of the contralateral visual hemifield (e.g., V1, MT). Other areas, however, introduce a discontinuity into the visual field representation at the horizontal meridian (e.g., the upper and lower field V2 representations are separated), lack a representation of part of the visual field (VP lacks a complete lower field map), distort or stretch their representation (e.g., V2, V3, VP, and VA and others are elongated in the central to peripheral direction), emphasize central vision or peripheral vision (inferotemporal and parietal areas, respectively), or have callosally supported representations of part of the

claustrum; and finally, layer I Cajal-Retzius neurons with widely branching tangential axons contacting pyramidal cell apical dendrites in layer 1. The local axon arbors of large pyramidal cell types are considerably larger than their dendritic fields; some smaller pyramids and stellate cells have more nearly comparable axonal and dendritic arbors.

There are many "standard" inhibitory types known or presumed to use GABA as a neurotransmitter that usually have smooth dendrites without spines and lack interareal axon branches entering the white matter (25, 26, 34, 35, 37, 38, 39). Some of these are: medium-sized layer 4C-alpha cells with 4C-alpha arbor and branch to layers 4A, 5A and 6; layer 6 cells back to layer 4C-alpha; layer 4C-beta cells to layers 4A and 3B; layer 4A cells back to layer 4C-beta; layer 4C-alpha cells with local axon spreading into layer 4B; layer 5A cells to layer 4C-alpha, 4A, 5A and 6; layer 5A cells with axon to only 5A and 3; basket cells in layers 3, 4, and 5 with beaded dendrites and wide-field local axons; medium-sized layer 4C-alpha cells with small field basket-like axons in layers 4C-alpha, 3B, 5A and 6; chandelier cells in layers 2, 3, 4, and 5 with small-field local axons terminating mostly on the spike-initiating part of pyramidal cell axons; double bouquet cells in layers 3 and 4 with extremely narrow vertically elongated local axon arbors (sometimes less than 20 microns wide) spreading across all layers; very small "clewed" cells with both dendrites and primary axon arbors restricted to layers 4C-alpha or 4C-beta; small cells with dendrites restricted to layer 4C-alpha or 4C-beta, but with axons spreading across both these layers; medium layer 4C-beta cells with dendrites spreading across layers 4C-beta and 4C-alpha but with axons restricted to 4C-beta. The axon arbors of inhibitory cell types may be similar size or much larger than their dendritic fields and in a number of cases, appear to provide points of contact between the cortical streams discussed above.

Finally, there are peptide-containing cells (40), presumed to have slow modulatory actions on post-synaptic cells. They have no clear axon and their vesicle-filled dendritic varicosities do not make conventional synapses. Some peptide-containing cells also appear to release GABA, a "standard" fast neurotransmitter. There are many neuropeptides (e.g., VIP, somatostatin, peptide Y, CCK, dynorphin, substance P) but it seems that there may be only two types of peptide cells--those containing VIP and those containing (simultaneously) several of the other peptides plus GABA. Most GABA cells do not contain peptides.

Most of the cell types described are present in other visual areas (42) and in somatosensory (34) and auditory areas. The main difference is probably that non-primary areas do not have as many cell types in layer 4. However, other areas probably have more cells types in other layers. Also, other areas may have cell types that have restricted tangential as opposed to laminar distributions (e.g., stripe-specific cells in V2).

2.4 Synapses

Our knowledge about the interconnections among different neuron types in the visual system at the level of the synapse is quite incomplete. Here I shall only make some general points about synaptic efficacy and connectional patterns based partly on more complete information available in other systems.

An excitatory or inhibitory cortical synapse normally involves a presynaptic "bouton" and a postsynaptic site on a cell body, on the flat surface of a dendrite, on a dendritic spine, or on the

initial segment of an axon. A single bouton often makes one synapse, but large boutons can make several. The axon of a single cell typically branches extensively to give rise to at least hundreds and perhaps as many as ten thousand boutons, and a target cell typically receives synapses from thousands of boutons. Recent studies in the spinal cord and brain stem suggest that both excitatory and inhibitory boutons behave in a probabilistic, all-or-none fashion, emitting either a single "quantum" of neurotransmitter or nothing at all when invaded by a spike (43, 44, 45, 46, 47, 48). (This contrasts with the better understood neuromuscular synapse, which typically releases several hundred quanta at once; however, the endplate is much larger than a bouton, and it drives a much larger target.) A single quantum produces a depolarizing (excitatory) or hyperpolarizing (inhibitory) postsynaptic potential of 100-200 microvolts in the target cell, measured at its cell body. Thus, a minimum of 50-100 of the excitatory boutons contacting a cell would have to be simultaneouly active (within about 5 milliseconds) to overcome the typical 10 millivolt threshold for generation of a spike. Models of the passive electrical properties of dendrites have suggested that the cell body potentials produced by distal dendritic synapses are attenuated and temporally broadened in comparision with potentials produced by otherwise equivalent proximal synapses. Consequently, distal synapses have traditionally been thought to be less efficacious in bringing a cell to threshold. However, the studies cited above suggest that distal synapses produce potentials at the cell body that are similar in size to those of proximal synapses, implying that distal synapses produce a larger local potential deflection. Since the distally-generated potentials have a longer time course at the cell body, they may in fact be quite efficacious.

These studies are significant because they allow us for the first time to attach rough numbers to anatomical pictures. Boutons are easily distinguished, for example, in HRP-filled axon arbors, even at the light microscopic level. A bouton distribution is difficult to assess, however, without knowledge of its postsynaptic targets. Fortunately, it has been possible in a few instances to inject both pre- and post-synaptic partners with HRP so that all contacts between the two can be visualized (43, 44, 45, 48, 49). The number of contacts between a pair of cells varies but has often turned out to be quite small. Connections that have been studied in this way include: Ia afferent contacts on motoneurons (one afferent makes 1-10 contacts per motoneuron--median 4); inhibitory cell synapses on Mauthner neurons (3-60 contacts--median 9); tectal output axons onto cells in the midbrain reticular formation (1-15 contacts--median 3); hair follicle afferents to spinocervical tract cells (1-60 contacts--a few for afferents supporting receptive field edges but many for afferents supporting center). Probably the most patent connection between cells is found in the cerebellum. Inferior olive axons branch to form several "climbing fibers" that each make hundreds of excitatory synapses onto single Purkinje cell dendrites. So far, a similar morphology has not been seen in the cerebral cortex.

Information about connections between neurons at the single cell level is scarce in the visual thalamus and cortex, and existing wiring diagrams have been constructed using mostly indirect data. Probably the best known connection involves the highly specialized chandelier cells or "axo-axonic" cells, which make 5-10 synapses with the axon initial segment of each of several hundred nearby pyramidal neurons (39). Up to 5 chandelier cells synapse on some pyramidal neurons. Thus, although a single chandelier cell might not be able to "veto" the output of a well-driven pyramidal cell (cf. ref. 1), a group of them probably could. Recently, the dLGN to V1

connection was examined at the single axon/single dendrite level, with some unexpected results (50). Single geniculate afferents provide only a few synapses to any given cortical neuron cell body or dendrite (estimated to amount to 1-8 contacts total per cell), but apparently contact very many of them (perhaps 400-800) in layer 4 of V1 in the cat. (Single monkey dLGN afferents in V1 have an order of magnitude fewer boutons than cat dLGN terminals; the number of contacts per cell they make is unknown).

Previously, it was thought that V1 cortical receptive fields were constructed from a much smaller number of dLGN cell inputs. It is perhaps not unexpected that a single axon in one of the highly non-topographic pathways from the tectum to the reticular formation (48) should be unable to single-handedly drive postsynaptic cells. The tectoreticular pathways mediate an activated-locus-to-firing-rate sensorimotor transformation that requires divergent/convergent connectivity in many parts of the reticular formation, and it may be impractical for one cell to make enough synapses. Surprisingly, however, low cell-to-cell connection strengths characterize even highly topographic projections like the one from dLGN to V1. In this case, single input axons make all their synapses in a small volume, but the synapses are not confined to few enough cells to single-handedly drive any of them.

In general, an axon in the cortex distributes boutons to a number of different cell types. Perhaps because of this, a cell of a given type does not usually make many connections with others like itself. Boutons from dLGN axons in layer 4 synapse on both excitatory and inhibitory cell types. They contact stellate and other non-pyramidal cells in layer 4 and pyramidal cells from other layers on spines and shafts and on basal and apical dendrites. Often the contact consists of a single bouton. Most inhibitory cells contacted by dLGN axons are in layer 4 because these cells tend to have smaller dendritic fields. The within-area targets of a few cortical cell types have been studied at the single-cell level, but the total number of synapses made by an axon on any given target cell is generally not known. Excitatory and inhibitory cortical axons, like dLGN axons, contact both excitatory and inhibitory cells, but the proportions and types of contacts vary. For example, the local collaterals of layer 6 pyramidal neurons contact the shafts of significant numbers of putative inhibitory cells in layer 4 (51), while the patchy wide-field local collaterals of large layer 3 pyramids make 90% of their contacts on layer 3 and 5 pyramidal neuron spines but only 5% of their contacts on inhibitory GABA cells (52). Inhibitory types also show quantitative differences. Chandelier cells only contact pyramidal types (39) (they may be unique in this), basket cells contact mostly excitatory types (pyramidal cell bodies, shafts and spines) but also some inhibitory types (26), and double bouquet cells contact mostly inhibitory types and only a few spines (38).

Since excitatory postsynaptic potentials and dendritic time constants are rather short, the typical firing rate of cortical input axons (usually under 100 Hz) is probably not fast enough to generate within-fiber temporal summation (50). Such a phenomenon may occur with inhibitory synapses since their potentials have longer time courses than those of excitatory synapses (20-40 milliseconds versus about 10 milliseconds). This asymmetry of time course may partly explain why there are fewer inhibitory cells and synapses than excitatory ones in the cortex. Since cells often receive 100 times the number of excitatory synapses needed to fire them, a given cell can probably be fired by many different input coalitions. Factors affecting the temporal coincidence

of different inputs may thus powerfully alter the effective connectivity of a local cortical network.

3 VISUAL SYSTEM PHYSIOLOGY

3.1 Excitatory Receptive Field Properties

The classically-defined excitatory receptive field of a cell is the region of the visual field from which a simple visual stimulus will elicit an excitatory response. The smallest receptive fields are found in V1 in the representation of the center of gaze and span less than a sixth of a degree of visual angle in macaque monkeys. Receptive field size increases with eccentricity in most areas. Receptive fields in secondary areas span several degrees, in tertiary areas, roughly five to ten degrees, in quaternary areas, over ten degress, and in upper level areas, even more. Even at the highest levels, however, few cells have excitatory receptive fields covering the entire visual hemifield. Receptive fields usually have sharply defined edges, even when they are quite large; almost invariably their borders are convex.

The original Hubel and Wiesel classification of orientation-tuned striate cortex cell types recognized three putatively serial stages (simple, complex, and hypercomplex cells), and has been reexamined in many studies from many different angles. Simple cells have spatially segregated subregions that respond to either the onset or offset of a small spot. Recent studies suggest that there are X-like (parvocellular) simple cells and Y-like (magnocellular) simple cells (53, 54). Some simple cell receptive fields have many "on" and "off" regions arranged as thin parallel strips (up to 8 total) so that they respond best when several cycles of a luminance grating are presented (54). Complex cell receptive fields do not have subregions when plotted with a spot; however, as with simple cells, some prefer more than one cycle of a grating (55). There may also be X-like and Y-like complex cells (53). It is known, now, that some complex cells are driven monosynaptically from the dLGN (56). The hypercomplex property (endstopping) is a graded property and may appear on the ends of both simple and complex cells. This property quite unexpectedly turned out to depend on the cortico-claustral loop (57). So far, few correlations between physiology and anatomy have emerged in cat V1. For example, there is no clear relation between receptive field size and dendritic field size and no obvious anatomical correlates of simple and complex cells. In fact, the one inhibitory cell type that has been characterized physiologically (the basket cell) can be either simple or complex, and be mono- or polysynaptically activated by either X- or Y-like inputs; physiologically, then, these cells are grossly indistinguishable from the various pyramidal cell types (26).

The responses of cells in a number of extrastriate visual areas have been tested with a variety of basic stimulus parameters (e.g., contrast, bar orientation, spatial frequency, direction of motion, speed of motion, wavelength, size, interocular disparity). Except for some simple cells in V2 (55), the receptive fields of cells in most extrastriate areas tend to be "complex" in the Hubel and Wiesel sense with respect to basic stimulus parameters, and most cells are binocular. In general, some cells in each area respond to changes in every parameter, although the proportion of cells responding to a particular parameter changes from area to area. In many cases, the sharpness of tuning for a parameter resembles that observed for this parameter in V1. Contrast affects the response of cells in all areas. All areas also contain cells sensitive to bar or grating

orientation, including the highest non-topographic areas in inferotemporal cortex (58). Most areas, again including inferotemporal cortex, contain cells tuned to spatial frequency (IT neurons often prefer many cycles within their receptive fields). The thick V2 cytochrome stripes, MT, V3 (DM), and their targets seem to be especially concerned with analyzing direction of motion (3, 4, 29, 30, 31). Cells in all areas have a broad tuning for stimulus speed (59). Isoluminant edges are specifically not detected by cells in MT and associated areas. Area DL shows many more cells tuned to the size of a stimulus than areas MT, DM, and M (the best stimulus size was often much smaller than the excitatory receptive field) (60). Cells in V2, V3, VP, and MT (at least) respond to horizontal (and in MT, also vertical) disparities (59, 61).

The sensitivity to several more complex stimulus parameters, and the effect of behavioral variables in awake animals has been examined, but usually only in one or two areas at a time. Some cells in V2 (but none in V1) respond to illusory contours generated by real contours just outside the excitatory receptive field (62). Some cells in MT are tuned to the overall direction of motion of a checkerboard pattern while other MT cells and all V1 cells are tuned only to the direction of movement of its component gratings (63). Some cells in area MST (but none in MT) respond to one or the other direction of rotation of a dot field, anywhere within their large receptive fields (64). Cells in MST and parietal cortex (but none in MT) respond to either contracting or dilating patterns (64, 65). Cells in MST respond during smooth pursuit eye movements, even when the target is temporarily extinguished (MT cells are silent without the target) (66). Attention to a simple spot stimulus, whether or not it is used as a target for a saccade, enhances the response to stimuli in much of the hemifield in parietal areas but not in inferotemporal areas (66, 67). Attention to one of two stimuli in a V4 or inferotemporal (but not a V1) receptive field suppresses the response to the other stimulus (68) (it is not clear whether this depends on attention to spatial location or to stimulus properties). Eye position signals can be used to sensitize different parts of the large receptive fields in area 7a (69) (it is not known if spatial attention can have such a localized effect in parietal areas). Finally, in a visual-visual or somatosensory-visual orientation match task, many cells in V4 increase their background firing and respond better to visual target stimuli after presentation of a cue to look for a particular orientation (different cells have different "cue preferences"), whether the initial oriented cue is presented visually or by touch (70). These results suggest that internally-generated activity in visual cortical areas (roughly, visual imagery) may appear quite early in the cortical areas hierarchy.

3.2 Receptive Field Surrounds

Silent inhibitory surrounds are well known at many levels in the visual system; they are "silent" since they become apparent only when the excitatory receptive field center is simultaneously stimulated. Typically, surrounds have been described as having an area comparable to that of the the excitatory center (see e.g., the hypercomplex end-zone). Recently, however, it has become clear that cortical neurons in all visual areas have much more extensive stimulus-specific inhibitory surrounds (71, 72, 73, 74, 75). In MT, random dots moving though the surround in the same direction as an optimally directed center stimulus can completely suppress the center response, while antagonistic surround movements enhance the response. The surround is typically 50-100 times the area of the excitatory center, and exerts its effect with about a 40 millisecond delay (71; cf. 53). Similar interactions have been demonstrated with respect to speed, color, spatial frequency, and orientation of center/surround pairs. This suggests that although most cortical neurons have local receptive fields, they do not signal purely local events except when given extremely sparse displays (like those in neurophysiology experiments).

3.3 Problems of Multiple Selectivities and Controls

The firing rate of a single neuron typically reflects selectivity to several different receptive field parmeters and to the nature of stimuli in the surround. It is intriguing to consider how subsequent processing stages might use such radically ambiguous information. A very large number of non-optimal combinations of center and surround parameters will give similar outputs. For instance, an MT neuron may significantly increase its firing rate if the direction or speed or contrast of either the center or the surround stimulus is changed. Furthermore, these parameters may interact nonlinearly in the neuron's response. From this perspective, it would be nice to have complete firing rate maps and also to know which parameters the neuron is actually insensitive to. Unfortunately, single neurons can rarely be held for the long periods (days—e.g., 5 factors by 5 levels each with 5 repetitions at 5 sec/repetition = 22 hours) such a map requires, and so such information is usually not available.

A related problem arises when more complex stimuli are employed, typically in higher areas. A reasonable strategy is to try to determine the 'simplest' explanation for why a cell is responding (e.g., orientation or spatial frequency content of a stimulus; simple spatial attention) and then invoke sensitivity to more complex stimulus configurations and behavioral states when simple explanations fail. As stimuli become more complex, the number of possible simple-explanation controls multiplies rapidly, and in most cases, we must be satisfied with a few spot checks.

4 STRATEGIES AND SUGGESTIONS FOR CONNECTING MODELS AND DATA

4.1 A Comparative Approach to the Data

In examining real nervous systems, it is important to realize that few structures are likely to exhibit their function transparently. Biological networks reflect the operation of many developmental and phylogenetic as well as purely functional constraints. Developmental programs are not infinitely plastic, but instead can be viewed as defining groups of possible network morphologies. Evolution does not produce the "best of all possible networks" for a particular task, but rather, ones that are just good enough, and ones that work within the developmental means of the brain region at hand. A five centimeter thick, 150 layer cortex may be ideal for certain visual tasks, but may not be developmentally practical. A corollary of this notion is that some aspects of the structure of a network may be non-functional—developmental excresences unavoidably generated in the process of making the functional parts, and persisting simply because they do little enough harm (cf. parasitic transistors in integrated circuit design). The notion that every structure or activity in an organism must have a function has been criticized before, though mostly with respect to non-neural morphology, as the "adaptationist fallacy" (76). None of this is to deny that some aspects of a neural circuitry represent remarkable functional adaptations.

To tease apart some of these influences, it sometimes helps to adopt an explicitly comparative perspective on function, comparing brain areas both within and between species. For example, the amygdaloid nuclei are not laminated like the cortex, and they have no vertically aligned elements like pyramidal cells. Nevertheless, the amygdala contains chandelier cells (77), which in the cortex, specifically inhibit the initial segments of pyramidal cells. Part of the explanation for why chandelier cells also appear in the amygdala (or conversely, in the cortex) is that the amygdala and the neocortex both develop from the pallium, the dorsal subdivision of the forebrain, and these two regions probably rely on similar mechanisms for determining cell types, and so on. The presence of these cells may not indicate that the "need" for the "chandelier" function in the amygdala (or the cortex) resulted in chandelier cells evolving there—another developmentally unrelated brain region might do it another way—but simply that the chandelier-generating machinery was handy. We thus might not want to provide a strictly functional explanation for their presence.

On the positive side, we can use parallel morphologies in unrelated brains or brain regions to attempt to argue that a particular function does in fact require a particular structure. One example concerns the tendency for evolutionarily derived ('advanced') brain regions to develop many layers and many cell types. The midbrain auditory nucleus torus semicircularis (=inferior colliculus) of a turtle, an animal with servicable though not acute hearing, is a small, unassuming structure with few layers and few cell types. By contrast, the same structure in gymnotid electric fish, which have a well-developed electric sense, has 12 layers and at least 50 cell types, and complex interlaminar connections (78) (the auditory and electrosensory sytems are both based on hair cell receptors and have closely related central pathways). Similarly, the cortex of 'primitive' mammals like insectivores and bats, but also dolphins (probably the first placental mammalian group to reenter the water) is thin and poorly laminated and appears (on quick inspection) to contain fewer and less differentiated cell types (79). The cortex of advanced mammals like seals, carnivores, and primates, by contrast, is much more clearly laminated and has many distinct cell types.

Other comparisons provide an interesting perspective on the function of the cortical streams discussed above. It has been suggested on the basis of finding double-opponent color-selective cells in the blobs of V1 in the macaque monkey that the blob system is specialized to process color. However, both owl monkeys and galagos have well-differentiated V1 blobs, in spite of their nocturnal habits and very poor color vision. Furthermore, the early primate ancestors of these groups were also probably nocturnal. Thus, it appears that color processing was added onto a previously existing system that did and may still mainly do something else (e.g., analysis of coarse shading) (80).

Clear contrasts, on the other hand, argue that a particular function identified in area A might not exist in area B. For example, it has been suggested that the cortex might perform computations based on fine-grained temporal interactions. When brain regions known to distinguish small time differences (less than 100 microseconds) are examined anatomically, however, they often show extremely specialized forms, including axons that synapse entirely on cell bodies (calyces of Held), turnip-like cells with extremely short dendrites, and electrical synapses, which all presumably function to preserve temporal information. The lack of such structures in the cortex suggests that operations with a similar temporal "grain" are not performed there.

4.2 Some Specific Suggestions

Subject to the cautions above, I list some general principles of organization in cortical visual systems. At the level of a single cortical area, there are many anatomical cell types (at least 50-100) with a great variety of local axonal arbor sizes and laminar specificities. Local collaterals of a cell may spread across several percent of an area. Within each area, there are many layers of cells and axon terminals (10 or more). In general, dendritic fields are an order of magnitude less widespread than local axon fields. Excitatory cell dendrites tend to cross more layers than inhibitory cell types. Most connections are between, not within cell classes. Cell to cell connections are weak (amounting to 1-5% of the firing threshold). Cells typically receive inputs from thousands of other cells. Excitatory and inhibitory types get both excitatory and inhibitory input (reflected in similar receptive field properties).

At the next level or organization, there is an interconnected hierarchy of distinct areas (20 or more) each characterized by a particular (probably quantitative) variation on the general pattern of within-area connectivity. Interareal connections are generally "acyclic" in the sense that there are (feedforward) paths to a given area of different lengths (i.e., that pass through different numbers of areas). Each area shows retinotopically-indexed excitation of increasing coarseness as one ascends, and wider field stimulus-specific inhibition. Between-area connections are strictly excitatory, usually topological (even at the highest levels), and they generate convex excitatory receptive fields at all levels. Each area contains cells with several different receptive field properties, but different areas have different mixes of cell properties. In general, lower areas in the hierarchy appear to have little "memory" (at least after development), reflected in their prompt, stable, non-habituating responses. Cells with very complex properties (e.g., "place cells" in the rat hippocampus that fire when the animal senses it is in a particular place in a familiar environment) may be only a few stages away from the input area (with place cells—V1, 18b, Ent, DG, CA3).

There are already instances where data begins to contact neural models more specifically. First, one class of abstract network models uses symmetric connectivity matricies (81), and the ubiquity of cortico-cortical feedback might at first seem to be an example of this. However, a round trip between areas requires considerable time (as much as 25 milliseconds if feedback influences have the same minimum latency as feedforward influences), and the return pathway is probably multisynaptic (some aspects of within-area connectivity may be more nearly symmetric). Such a feedback delay may in fact lead to interesting temporal pattern sequences (82) with the persistence time of each pattern equal to the feedback delay. So far, such transition and hold patterns have not been obvious in recordings from early, stimulus-bound levels in the areas hierarchy, although no one has explicitly looked for them. Given typical cortical neuron firing rates, each pattern could be represented by only a few spikes in a single neuron. Second, a number of models result in neural-like units being driven to saturation (or silence) as the network approaches a stable state. This is not usually seen in the cortex (cells tend to show tuning curves), but as before, it may occur in less well studied higher areas, where the cells

sometimes tend to burst and habituate. Third, there are a variety of models of how motion information is used to infer structure and other things. The experiments with translation, dilation, and rotation of dot fields in parietal cortex suggest that this is done in at least two stages, with the first stage (MT) sensitive only to translation and the next stage (MST, 7a, and others) sensitive to all three. This data argues against one stage models. Finally, there has not yet appeared evidence of explicit z-axis information at the single cell level in any cortical area, in great contrast to the many explicit representations of variously blurred x- and y-axis information. This argues that models with explicit z-axis coordinates are too abstract to make contact with cortical mechanisms.

These notes are very brief, but they may suggest that a good plan might be to use known neuron properties as starting point. A particular cortical area doesn't "know" where, say, the direction selectivity or end-stopping in its input came from, and is thus a bit like a neural modeler. By using what is already known at one level in the cortical hierarchy, it may be possible to suggest plausible response properties at the next level that neurophysiologists could search for in the brain. It seems unlikely that a blind empirical approach to higher areas is going to produce major insights by itself.

4.3 Other Analogies

A class of network models has been shown to be formally similar to physical systems known as spin glasses, of which a cool magnet is an example. Analogies with other physical systems may be possible. For instance, there is an interesting code-based physical system in living cells that controls a great variety of chemical reactions and distributions in a thoroughly 'parallel' fashion (cell metabolism) through the use of high-specificity reaction-controlling devices (enzymes and other proteins). Although spin glasses may have a "rich statistical physics" (83), there are no living magnets. Living physical systems are predicated on making and breaking strong covalent bonds in addition to exploiting weak cooperative interactions like those seen in spin glasses. This suggests that we might do well to look for analogues to strong (covalent) bonding and bond-making and breaking processes between stable 'molecular' firing patterns (perhaps embodied in strongly interacting cell subassemblies) in neural networks. The metabolic network in code-based living chemical systems, of course, depends heavily on objects whose detailed three-dimensional configurations are determined by cooperative relaxation (protein folding). But protein self-folding occurs only after the arena of possible weak interactions has been defined by the serial synthesis of a 'one-dimensional' chain of peptides using strong covalent bonds.

It may be possible to build up an analogous 'metabolism' of patterns in a network controlled by other serially assembled and then 'folded' patterns. In any case, it will be necessary to go beyond the simple stimulus-response framework that characterizes most present day network models (and neurophysiology experiments). Eventually, we need to approach questions of how, for instance, one integrates successive glances in scene perception or successive words in language comprehension, and how the integrated scene or discourse representations interact appropriately with their previously constructed fellows in the brain.

5 CONCLUSION

A great deal of data had to be very briefly summarized here. Even so, it is clear that our

knowledge of many basic parameters in the visual system is rudimentary. And the suggestions given on what to do with some of this information are, to be polite, preliminary. But it is emphatically necessary to start somewhere. Neural modelers are fond of pointing to the futility of trying to determine how a large computer functions by taking 'single-unit' recordings from it (81). I think such pronouncements underestimate human experimental and theoretical ingenuity (and luck). We will get better theories and data by trying to connect the two. Some of the intermediate stages in this development are bound to turn into curiosities known only to future historians of science. We shall not, however, be able do without intermediate stages.

6 ACKNOWLEDGEMENTS

I thank T. Sejnowski for helpful discussions and C. Hochenedel for preparing the manuscript. Supported by NIH grant F32 EY05887.

7 REFERENCES

(1) F. Crick, C. Asanuma, In Parallel Distributed Processing, Volume 2: Psychological and Biological Models, MIT Press, Cambridge 1986, pp. 333-371.

(2) G.A. Orban, Neuronal Operations in the Visual Cortex, Springer-Verlag, Berlin 1984.

(3) D.C. Van Essen, In Cerebral Cortex, Volume 3: Visual Cortex, Plemun Press, New York 1985, pp. 259-329.

(4) J.H.R. Maunsell, W.T. Newsome, Ann. Rev. Neurosci. 10 (1987) 363-402.

(5) J.H. Kaas, Ann. Rev. Psych. 38 (1987).

(6) J.F. Baker, S.E. Petersen, S.E. Newsome, J.M. Allman, J. Neurophysiol. 45 (1981) 397-416.

(7) R.E. Weller, J.H. Kaas, J. Comp. Neurol., 256 (1987) 137-172. See references within.

(8) M.I. Sereno, C.T. McDonald, J.M. Allman, Neurosci. Abstr. 12 (1986) 1181.

(9) D. Felleman, J.J. Knierim, D.C. Van Essen, Neurosci. Abstr. 12 (1986) 1182.

(10) M.M. Merzenich, J.H. Kaas, Prog. Psychobiol. Physiol. Psych. 9 (1980) 1-42.

(11) D.N. Pandya, E.H. Yeterian, In Cerebral Cortex, Volume 4: Association and Auditory Corticies, Plenum Press, New York 1985, pp. 3-61.

(12) R.W. Rodieck, In Comparative Primate Biology, Volume III, Neurosciences, Alan R. Liss 1987 (in press).

(13) C.-S. Lin, J.H. Kaas, J. Comp. Neurol. 187 (1979) 655-678.

(14) J. Graham, C.-S. Lin, J.H. Kaas, J. Comp. Neurol. 187 (1979) 557-580.

(15) H.R. Rodman et al., Neurosci. Abstr. 11 (1985) 1246.

(16) C. Asanuma, R.A. Andersen, W.M. Cowan, J. Comp. Neurol. 241 (1985) 357-381.

(17) J.H.R. Maunsell, In Matters of Intelligence, Reidel Press, Dordrecht 1987 (in press).

(18) S.E. Petersen, F. Miezin, J. Allman (1987, in press).

(19) B.J. Richmond, R.H. Wurtz, T. Sato, J. Neurophysiol. 50 (1983) 1415-1432.

(20) D.G. Amaral, In Handbook of Physiology, Williams and Wilkins 1987 (in press).

(21) J.S. Lund, J. Comp. Neurol. 147 (1973) 455-496.

(22) J.S. Lund, R.G. Boothe, J. Comp. Neurol. 159 (1975).

(23) D. Fitzpatrick, J.S. Lund, G.S. Blasdel, J. Neurosci. 5 (1985) 3329-3349.

(24) G. Blasdel, J.S. Lund, D. Fitzpatrick, J. Neurosci. 5 (1985) 3350-3369.

(25) F. Valverde, In Cerebral Cortex, Volume 3: Visual Cortex, Plenum Press, New York 1985, pp. 207-257.

(26) K.A.C. Martin, In Cerebral Cortex, Volume 2: Functional Properties of Cortical Cells, Plenum Press, New York 1984, pp. 241-284.

(27) R.B.H. Tootell, M.S. Silverman, R.L. DeValois & G.H. Jacobs, Science 220 (1983) 737-739.

(28) M.S. Livingstone, D.H. Hubel, J. Neurosci. 4 (1984) 309-356.

(29) S. Shipp, S. Zeki, Nature 315 (1985) 322-325.

(30) D.H. Hubel, M.S. Livingstone, Nature, 315 325-327.

(31) E.A. DeYoe, D.C. Van Essen, Nature 317 58-61.

(32) C. Enroth-Cugell, J.G. Robson, J. Physiol. 187 (1966) 517-552.

(33) Y. Fukada, J. Stone, J. Neurophysiol. 37 (1974) 749-772.

(34) E.G. Jones, J. Comp. Neurol., 160 (1975) 205-268.

(35) T.L. Davis, P. Sterling, In Models of Visual Cortex, John Wiley, New York 1985, pp. 524-532.

(36) C.D. Gilbert, T.N. Wiesel, J. Neurosci. 3 (1983) 1116-1133.

(37) J.S. Lund, J. Comp. Neur. 257 (1987) 60-92.

(38) P. Somogyi, A. Cowey, In Cerebral Cortex, Volume 1: Cellular Components of the Cerebral Cortex, Plenum Press, New York 1984, pp. 337-360.

(39) A. Peters, In Cerebral Cortex, Volume 1: Cellular Components of the Cerebral Cortex, Plenum Press, New York 1984, pp. 361-380.

(40) E.G. Jones, S.H.C. Hendry, Trends Neurosci. 9 (1986) 71-76.

(41) A.J. Rockel, R.W. Hiorns, T.P.S. Powell, Brain 103 (1980) 221-244.

(42) J.S. Lund, A.E. Hendrickson, M.P. Ogren, E.A. Tobin, J. Comp. Neurol. 202 (1981) 19-46.

(43) J.J. Jack, S.J. Redman, K. Wong, J. Physiol. 321 (1981) 65-96.

(44) S. Redman, B. Walmsley, J. Physiol. 343 (1983) 117-133.

(45) S.Redman, B. Walmsley, J. Physiol. 343 (1983) 135-145.

(46) E.A. Neale, P.G. Nelson, R.L. Macdonald, C.N. Christian, L.M. Bowers, J. Neurophysiol. 49 (1983) 1459-1468.

(47) H. Korn, A. Mallet, A. Triller, D.S. Faber, J. Neurophysiol. 48 (1982) 679-707.

(48) M.I. Sereno, J. Comp. Neurol. 233 (1985) 48-90.

(49) A.G. Brown, Organization in the Spinal Cord, Plenum Press, New York 1981.

(50) T.F. Freund, K.A.C. Martin, P. Somogyi, D. Whitteridge, J. Comp. Neurol. 242 (1985) 275-291.

(51) B.A. McGuire, J.-P. Hornung, C.D. Gilbert, T.N. Wiesel, J. Neurosci. 4 (1984) 3021-3033.

(52) Z.F. Kisvarday, K.A.C. Martin, T.F. Freund, Z. Magloczky, D. Whitteridge, P. Somogyi, Exp. Brain Res. 64 (1986) 541-552.

(53) J. Bullier, M.J. Mustari, G.H. Henry, J. Neurophysiol. 47 (1982) 417-438.

(54) L.A. Palmer, J.P. Jones, W.H. Mullikin, In Models of the Visual Cortex, Wiley, New York 1985, pp. 273-280.

(55) D.A. Pollen, K.H. Foster, J.P. Gaska, In Models of the Visual Cortex, Wiley, New York 1985, pp. 281-291.

(56) G.H. Henry, In Cerebral Cortex, Volume 3: Visual Cortex, Plenum Press, New York 1985, pp. 119-155.

(57) H. Sherk, S. LeVay, J. Neurosci. 3 (1983) 2121-2127.

(58) D.A. Pollen, M. Nagler, J. Daugman, R. Fronauer, P. Cavanagh, Vision Res. 24 (1984) 233-241.

(59) A. Burkhalter, D.C. Van Essen, J. Neurosci. 6 (1986) 2327-2351.

(60) S.E. Petersen, J. Baker, J.M. Allman, Brain Res. 197 (1980) 507-511.

(61) J.H.R. Maunsell, D.E. Van Essen, J. Neurophysiol. 49 (1983) 1148-1167.

(62) R. von der Heydt, E. Peterhans, G. Baumgartner, Science 224 (1984) 1260-1262.

(63) J.A. Movshon, E.H. Adelson, M.S. Gizzi, W.T. Newsome, Exp. Brain Res. (1987, in press).

(64) H. Saito, M. Yukie, K. Tanaka, K. Hikosaka, Y. Fukada, E. Iwai, J. Neurosci. 6 (1986) 145-157.

(65) M.A. Steinmetz, B.C. Motter, C.J. Duffy, J. Neurosci. 7 (1987) 177-191.

(66) R.H. Wurtz, B.J. Richmond, W.T. Newsome, In Dynamic Aspects of Neocortical Function, Wiley, New York 1984, pp. 195-217.

(67) S.E. Petersen, D.L. Robinson, W. Keys, J. Neurophysiol. 54 (1985) 867-886.

(68) J. Moran, R. Desimone, Science 229 (1985) 782-784.

(69) R. Andersen, G. Essick, R. Siegel, Science 230 (1985) 456-458.

(70) P.E. Haenny, J.H.R. Maunsell, P.H. Schiller (1987, in press).

(71) J. Allman, F. Miezin, E. McGuinness, Perception 14 (1985) 105-126.

(72) J. Allman, F. Miezin, E. McGuinness, Ann. Rev. Neurosci. 8 (1985) 407-430.

(73) J.I. Nelson, B. Frost, Brain Res. 139 (1978) 359-365.

(74) S. Zeki, Neurosci. 9 (1983) 767-781.

(75) J. Moran, R. Desimone, S.J. Schein, M. Mishkin, Neurosci. Abstr. 9 (1983) 957.

(76) S.J. Gould, R.C. Lewontin, Proc. R. Soc. Lond. B 205 (1979) 581-598.

(77) A.J. McDonald, J. Comp. Neurol. 212 (1982) 293-312.

(78) C.E. Carr, L. Maler, J. Comp. Neurol. 235 (1985) 207-240.

(79) P.J. Morgane, M.S. Jacobs, A. Galaburda, In Dolphin Cognition and Behavior, Lawrence Erlbaum, Hillsdale, NJ 1986, pp. 5-29.

(80) J. Allman, M. Sereno, Dahlem conference (in press).

(81) J.J. Hopfield, D.W. Tank, Science 233 (1986) 625-633.

(82) H. Sompolinsky, I. Kanter, Phys. Rev. Lett. 57 (1986) 2861-2864.

(83) J.J. Hopfield, D.W. Tank, Science 235 (1987) 1228-1229.

Design for a Sensorium

Jan J. Koenderink

SUMMARY

Two problems in brain theory as related to the structure of sensoria are discussed: how to construct a topology of the visual field and how to embed differential geometry in the representation in order to enable the construction of shape measures. Possible embodiments in terms of the known physiological structure are investigated.

A topology can be based on the correlation structure of neural activity. Under certain constraints this can be shown to lead to a neat triangulation, i.e. the dimension and the connectivity of the space are obtained from scratch.

Sets of receptive fields (RF's) can act as local "n-jet's": truncated Taylor series that locally approximate the retinal illuminance distribution. The RF's can be considered "fuzzy partial derivatives". In such a view many apparently multilocal properties (such as boundary curvature) can be considered simple point properties, thus diminishing the level of geometrical expertise required from sensorimotor routines addressing the sensorium.