MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 1311 December 1991

Sequence-Seeking and Counter Streams: A Model for Information Processing in the Cortex

Shimon Ullman

Abstract

This paper presents a model for the general flow of processing in the neocortex. The basic process, called "sequence-seeking," is a search for a sequence of mappings, or transformations, linking source and target representations. The search is bi-directional, "bottom-up" as well as "top-down," and it explores in parallel a large number of alternative sequences. This operation is implemented in a structure termed "counter streams," in which multiple sequences are explored along two separate, complementary pathways which seeking to meet. The first part of the paper discusses the general sequence-seeking scheme and a number of related processes, such as the learning of successful sequences, context effects, and the use of "express lines" and partial matches. The second part discusses biological implications of the model in terms of connections within and between cortical areas. The model is compared with existing data, and a number of new predictions are proposed.

© Massachusetts Institute of Technology, 1991

This paper describes research done at the Department of Brain & Cognitive Sciences and the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. This research is sponsored by grants from the National Science Foundation under contract IRI-8900207, and the Air Force Office of Scientific Research under contract 89-504. Support for the A.I. Laboratory's artificial intelligence research is provided by ONR contract N00014–91–J–4038. Shimon Ullman is affiliated with MIT and with the Department of Applied Mathematics & Computer Science at the Weizmann Institute of Science, Rehovot 76100, Israel.

THE UNIVERSAL OPERATION OF THE CORTEX

This paper presents a hypothesis regarding the basic scheme of information processing in the neocortex. The neocortex (the newer part of the cortex) is a relatively recent part of the brain, responsible for may of its advanced functions. In humans, it accounts according to some estimates for as much as 70% of the cells in the central nervous system (Nauta & Feirtag 1986), and it is involved in a variety of functions, including sight, hearing, somatosensory perception, motor planning, language, and general aspects of thinking, problem solving, and planning.

Are there some general principles of information processing that are used, with suitable variations, to carry out different cortical functions? Our present state of knowledge does not allow an unequivocal answer to this fundamental question, but the notion of a common, basic, cortical scheme has a long history, and has been discussed recently. e.g., by Creutzfeldt (1978), Barlow (1985), Martin (1988), and White (1989). A number of researchers and Poggio (1990) have proposed different hypotheses at different levels of detail regarding the basic cortical scheme (some of these proposals are reviewed briefly in the next section).

The main support for the notion of a "universal" cortical scheme comes from the similarity in structure between different cortical areas (e.g., visual, auditory, sensory-motor, etc.), and even cortical structures in different animals. As expressed by Martin (1988, p. 639), "...it would take an expert to distinguish rat frontal cortex from sheep parietal cortex, or cat auditory cortex from monkey somatosensory cortex." Cell numbers and the distribution of cell types also appear to be similar across species and in different cortical regions. For example, Rockel *et. al.* (1980) have compared cell counts throughout the thickness of the neocortex in different animals, and found that, with the exception of area 17 in a number of primates, cell counts in different functional areas (motor, somatic sensory, area 17, frontal, parietal and temporal), and in several species (mouse, cat, monkey, man) were almost identical.

There are, of course, cytoarchitectonic differences among cortical areas, for example, the striate cortex in certain primates, or the giant Betz cells in the motor cortex. In fact a view expressed by many neuroanatomists

in the earlier part of this century (Broadman, Meynert, Voght and Voght, von Economo, and others, see, e.g., in Creutzfeldt (1978)) suggested that the cytoarchitectonic differences among cortical areas were crucial to the specific functions of the different areas.

With later studies, as new techniques started to reveal in more detail the afferent and efferent structures, as well as the internal connectivity in the cortex, the view shifted more towards an emphasis on the similarities in structures in different parts of the neocortex. A strong version of this view was expressed by Creutzfeldt (1978), who suggested that "the specific functional and morphological differences of the various cortical areas can be considered unessential. The 'specific function' of any cortical areas are defined by the origin of its afferents and the destination of its efferent connections" (p. 377). Barlow (1985) expressed a similar view: "...regional differences of structure are interesting and important but they do not violate the hypothesis that the cortex performs a uniform operation throughout" (p. 38), and, in the same spirit, White (1989), in his review of cortical circuits, raises the notion of "some common, basic plan" of the cerebral cortex.

Such a common, basic plan, using common machinery, may be advantageous from the point of view of the evolution of the neocortex. It will make it easier to expand the capacity of the cortex by replicating existing cortical areas and recruiting the new areas for new functions.

It should be noted, however, that the uniformity of structure and function can be expected to hold in general outline only. As noted by Crick (1986): "...it is reasonable to assume that the neocortex has evolved because it is good at a particular sort of computation that, with appropriate local variations, it may be broadly the same in all parts. We may also expect that these basic processes will be complicated by somewhat elaborate neural gadgetry designed to make for better performance" (p. 369).

In addition to the structural similarities, some additional support for the notion of a common cortical scheme comes from the work of Sur and his collaborators (Sur, Garraghty and Roe 1988, Pallas, Roe, and Sur, 1990, Roe, Hahm and Sur 1991) on re-directed visual input. In this work, retinal output in the ferret was induced to project to the auditory thalamus (MGN) that projects in turn to the primary auditory cortex (A1). An ordered retinotopic map was found to be established in A1 as a result of its novel

visual input. Furthermore, receptive fields in A1 in rewired ferrets were found to be strikingly similar to those recorded in primary visual cortex of normal ferrets (Roe, 1991). These findings are consistent with the general notion of schemes of information processing shared by different areas of the neocortex.

In his review of circuits in the cerebral cortex, Martin (1988) suggests that "..the tempting notion is then that nature's laboratory has hit on a process that enables it to use the same machinery for very different ends" (p. 640). The goal of this paper is to make some speculative, but concrete proposals regarding this common machinery.

The Cortical Scheme: Previous Approaches

Several proposals have been made in the past regarding the basic operation of the neocortex. In this section I briefly list seven of these proposals.

(i) *The classifying cortex.* In Marr's (1970) theory, the basic function of the neocortex is the classification of incoming patterns. New classes are formed by an ongoing process of clustering frequent and closely similar patterns. Incoming stimuli are then tested for membership in already existing classes. The classification is performed in two stages. In the first, the incoming pattern is represented in terms of so-called "codon" cells that "look for" specific sub-patterns in the incoming stimulus. In the second stage, "classificatory" units, representing different classes of events, are driven by the codon cells. The theory identifies the codon cells primarily with the Martinotti cells, a class of large layer 6 pyramidal cells, and the output classificatory units with layer 3 and 5 pyramidal cells. Inhibitory interneurons are assigned two main tasks: controlling the dynamic range of the output units (by subtractive inhibition), and setting their threshold (by divisive inhibition).

(ii) *Nonlinear spatio-temporal filter.* Creutzfeldt (1978) proposes a view of the cortical network as a nonlinear spatio-temporal filter linking the activity in the thalamus and other afferent structures to the effectors. This filtering operation allows through and emphasizes only certain aspects of the signals transmitted to the cortex. In particular, it emphasizes change, movement, and repetitive stimulation. In this scheme the uniformity of the

cortex reflects a uniformity in the filtering operation applied to different afferents.

(iii) *Coincidence detection.* In Barlow's (1985) scheme, the common function of units throughout the neocortex is "the detection and signalling of suspicious coincidence." A coincidence between input events A and B is "suspicious" if $P(A\&B) >> P(A) \times P(B)$. The general suggestion is that such coincidences are used by the cortex to form internal models of the environment. Hypotheses made by the cortex about the environment imply certain coincidences among different inputs. Such coincidences are in turn detected and used to support or refute the hypotheses. A simple Hebbian mechanism is proposed for constructing "coincidence detectors": synaptic connections to a given unit increase in efficacy for inputs that co-occur frequently.

(iv) *Multi-level relaxation.* In Mumford's (1990) theory, multiple cortical areas, each representing restricted aspects of the world, interact to achieve a consistent interpretation of the incoming stimulus. The areas are organized hierarchically, with low-level areas concerned with aspects that are closer to the raw data (e.g., grouping together line segments), and higher-level areas concerned with more abstract aspects (e.g., identifying a face). The interpretation process involves a relaxation process that uses two paths of information flow: a cortical one, and a thalamo-cortical one. The cortical path involves interactions between reciprocally connected cortical areas. In this interaction, the higher, more abstract area projects to the lower area a "template," an instance of a model stored in the higher area (e.g., the shape of a circle of a particular size and location). The signal from the lower to the higher area carries a residual: the remaining differences between the data in the lower area and the "hypothesis" of the higher one. The relaxation process proceeds until the residual becomes insignificant or can be accounted for, or else the hypothesis is replaced by an alternative one. In the thalamo-cortical loop, subdivisions of the thalamus serve as so-called "active blackboards" to summarize and integrate the information from the cortical areas. (This description applies to the sensory part of the cortex; a related part of the theory deals with actions and plans). A somewhat similar proposal is the "adaptive resonance theory" of Carpenter

and Grossberg (1987, Grossberg 1988) in which patterns are classified by a combination bottom-up and top-down processes.

(v) *Large-scale associative memory.* A number of models have proposed to view the cortex primarily as a large-scale associative memory (or a collection of associative memories). An associative memory is a device in which a large number of patterns (P_i) can be stored. Given an input pattern, the system can retrieve, from all the stored patterns, that pattern P_i which most closely resembles the input. This memory has the power of error-correction and completion; the input pattern may be a corrupted and incomplete version of P_i, and yet the correct P_i can be retrieved. Willshaw et al. (1969), Hopfield (1982), and Amit (1989), among others, have described possible neural implementations of associative memories, as well as possible extensions of the basic memories to solve a number of problems.

(vi) *Interpolating memory.* In Poggio's (1990) HyperBF scheme, the brain begins the learning of a new task by acquiring a sufficient number of input-output examples. For a new, unknown input, the appropriate output is then derived by an interpolation (or approximation) process applied to the stored data. For example, a number of motor programs, to reach different final configurations, may be stored in memory. A new program, to reach a new goal, can then be synthesized by approximation in the space of motor programs. (A related scheme was proposed by Albus (1971) for the cerebellum.) The approximation modules in the brain are proposed to be "hyperBF modules," a generalization of the Radial Basis Functions technique used in approximation theory. In this technique, an output value is obtained by the superposition of basis functions stored in memory and centered around previously learned examples. In a biological implementation, the basis functions are proposed to be multi-dimensional Gaussians, which can then be decomposed into the product of low-dimensional ones.

(vii) *Theory of neuronal group selection.* This theory and its implementation in large scale computer simulations have been developed by G. Edelman and his collaborators (e.g., Edelman 1978, Reeke *et al.* 1990).

In this theory the basic functional unit of the cortex is a group of hundreds to thousands of strongly interconnected neurons. These basic groups,

or 'repertoires', then form higher order structures such as maps and pathways. Maps along a pathway and between parallel pathways are interconnected by rich reciprocal 're-entrant' connections.

Local rules of synaptic modification can constantly modify the connections between groups, as well as the structure of the underlying groups. The process driving these changes is likened to an evolutionary selection process. Groups and connections that lead to useful behavioral consequences will be favored, leading over time to more adaptive behavior.

The theory has been applied in simulations to various problems, including the organization of maps in somatosensory cortex, the integration of multiple visual maps, and aspects of pattern recognition and visual-motor integration.

The current model: Sequence-seeking in counter streams structure Briefly, the basic process in the model, called "sequence-seeking", is a search for a sequence of mappings, or transformations, linking source and target representations. The search is bi-directional, "bottom-up" as well as"top-down", exploring simultaneously a large number of alternatives. This operation is implemented in a structure termed "counter streams", in which multiple sequences are explored along two separate, complementary pathways, seeking to meet. Roughly speaking, top-down information, originating from stored models, flows along a descending path, and bottom-up information, originating from the sensory data (in the case of perception) flows up along an ascending stream. When similar patterns are reached from both ends, an entire link connecting the source and target representations is formed, and also learned for future use.

As we shall see, the model proposed in this paper is not incompatible with some of the previous proposals. In the current model, the uniform operation of the cortex has to do primarily with establishing sequences of mappings, or transformations, and it is not directly concerned with the individual mappings themselves. Much of the processing is undoubtedly accomplished within the individual mappings, perhaps using some of the mechanisms described by the models listed above. The model described below is not concerned, however, with the computations performed by specific mappings in specific cortical areas, but with possible common principles that do not depend upon the details of individual mappings.

Main Properties of the Model

Part I of the paper describes the proposed scheme, the sequence-seeking model, and the counter-streams structure that implements it. As an introduction, the main properties of this model are also summarized briefly in the list below. The list merely introduces the main notions and the order of presentation, without explanation. The issues are then discussed in detail in Part I. Part II discusses the biological implications of the model, compares it with existing data, and proposes new predictions.

- The basic process, termed 'sequence-seeking', is a search for a sequence of mappings, or transformations, linking a source and target representations.

- The search is bi-directional, 'bottom-up' as well as 'top-down'.

- The two searches are performed along two separate, complementary pathways, an ascending one and a descending one, seeking to meet (the 'counter streams' structure). In the cortex, these two streams use the forward and backward-type connections, respectively.

- There are cross connections of the 'priming' type between the two streams.

- A track being traversed in one stream leaves behind a primed trace in the complementary stream.

- During the search, many alternative sequences are explored in parallel, using a maximal possible number of co-active patterns. Active patterns persist for a while, then they decay and are replaced by new ones.

- When two sub-sequences, one from 'above', the other from 'below', meet, they generate a chain of activity linking the source and target representations.

- The ascending and descending streams can meet at any level.

- The entire sequence that has been used successfully is increased in

relative strength. This is obtained by a local learning rule, based on the approximate temporal coincidence of activation in the two streams.

- Context effects are mediated by partial tracks being primed, for longer time intervals.

- 'Express lines' can traverse more than a single level, to promote the expansion of favored sequences.

- Favorable sequences, determined by the degree of match between the ascending and descending streams, will be expanded further preferentially, at the expense of less promising sequences.

- A partially successful match between patterns in the ascending and descending streams is used to constrain further transformations of this pattern.

- Connections of different types (signifying different types of relations between patterns) are implemented (a) by using different sub-structures in the network, and (b) by using associations between subpatterns.

- An additional loop serves to control the process, e.g. by selecting restricted regions in the visual field, and by controlling the use of different sub-structures participating in the computation. In the cortex, this loop may involve layer 5 and subcortical structures.

- Most of the processing is performed by the expanding sequences, rather than by a prolonged relaxation or iterative computation within a single area.

PART I
The Sequence-Seeking Scheme and Its Implementation in the Counter-Streams Structure

The goal of this part is to describe the proposed scheme of basic cortical operation, the sequence-seeking scheme, and the counter-streams structure that implements it. The discussion in this part is in computational, or information-processing terms; the next part discusses in more detail the

proposed biological implementation. The sequence-seeking scheme is presented primarily in the context of perception, but the proposal is intended to apply in general terms to other cortical functions as well.

Applying sequences of transformations

In perception, the brain is faced with the task of establishing a link between an incoming pattern and stored representations. The two will often fail to match exactly; for example, one may see a shape that differs in position and size from previously stored representations of the same shape. The differences can be reconciled by applying a series of transformations to the shape, such as scale and position changes. A fundamental aspect of the basic cortical operation, according to the sequence-seeking scheme described below, is that it seeks to establish a link between different representations or patterns, by exploring sequences of transformations.

The search for a connecting sequence of transformations, mappings, or states, linking source and target representations, could provide a useful mechanism for various aspects of perception, as well as for other, non-perceptual functions. For example, the planning of a motor action can be cast at some level in terms of seeking a sequence of possible moves, linking an initial configuration with a desired final state. Visual routines (Ullman 1984), which are 'programs' for extracting shape properties and spatial relations (such as "A is inside B") have an analogous structure. Similarly, more general planning and problem solving can often also be described in terms of establishing a sequence of transformations, mappings, or intermediate states, linking some source and target representations.

Bi-directional search

The sequence-seeking is bi-directional, "bottom up" as well as "top down." In applying a sequence of transformations to match an incoming pattern P with a stored pattern M, the transformations could be applied to P, or to M, or to both. In general, the only viable alternative is to apply the transformations to both the incoming and stored patterns. The reasons for the bi-directional search are the following. First, some of the transformations can only be reasonably applied to the stored pattern M. Suppose that M represents, for instance, a pair of scissors. There are certain transformations that scissors can undergo (opening and closing) that other

objects cannot. These transformations can therefore be applied to the model M but not to the image P. More universal transformations such as scaling and rotation can and should be applied to P rather than to M. The reason is that before we know what the object is, we attempt to match it with a number of different stored representations of different objects. Applying the transformations to all of them will lead to a proliferation of patterns that can be avoided by transforming P alone.

In addition, splitting the search for a sequence of transformations into two searches from opposite ends often increases the efficiency of the computation, and reduces the overall number of patterns being explored. These considerations suggest the use of two streams of transformation, an ascending one (starting at the stored models) and a descending one (starting at the image), seeking to meet. As we shall see, the suggestion is that from a biological standpoint, these streams correspond to the "forward" and "backward" connections between cortical areas.

Exploring multiple alternatives

A large number of alternative routes may have to be explored before a link is successfully established between a "source" and a "target" representation. To achieve fast computation, it will be necessary to explore simultaneously a large number of alternative routes.

For a number of reasons, it will not be possible to use the extreme alternative of exploring all of the possible sequences at the same time. First, the number of simultaneously co-active patterns will become prohibitively large. Second, in the proposed neural implementation (discussed in Part II), patterns are represented by a distributed activity over a population of neurons. Since different patterns are represented by the same group of neurons, such a scheme places limitations on the number of patterns that can be co-active without leading to confusion (see Appendix). The system should therefore be wired up in such as way that at any given moment a maximal number of alternatives will be explored in parallel, subject to the limitations placed on the co-activation of patterns. Patterns will become active for a certain time period, then they will decay, and be replaced by alternative patterns that have not yet been explored.

In many models of visual processing, the input pattern undergoes a single sequence of processing stages. In contrast, one of the distinguishing

characteristics of the sequence-seeking scheme is that an input pattern gives rise to multiple sequences of transformations and mappings that are being explored in parallel. The terms "transformations" and "mappings" should be taken here in a broad sense; they may include geometric transformations such as changes in size, position, and orientation, the recovery of different properties such as color, motion, texture, and 3-D shape, as well as exploring alternative ways of representing the pattern in terms of its parts and its abstract shape properties.

In exploring multiple sequences, it is also possible that a source and a target representation will be linked by more than a single path. This may be in accord with the evidence that different aspects of the stimulus may be analyzed by the visual cortex along separate (though interacting) pathways (e.g., Zeki & Shipp 1988). This possibility is compatible with the sequence-seeking scheme, but will not be discussed here in detail.

The existence of multiple alternative sequences will also increase the robustness of the system. For example, if a lesion to the system disrupts some established sequences, additional experience could lead to the establishment of new, perhaps less direct sequences. This is consistent with many findings regarding the recovery from visual deficits caused by lesions (e.g. to area V4, Desimone *et al.* 1990, Schiller & Lee 1991). Training following the lesion usually results in considerable recovery in performance level. The recovery is task-specific, and, although performance improves, there are often lasting effects in terms of increased response time.

Linking the ascending and descending streams: the counter-streams structure

The scheme that emerges so far is summarized schematically in Fig. 1a. The basic operation of the cortex in this scheme is to seek a sequence of transformations linking (in the case of perception) an incoming pattern (S in Fig. 1a) with a number of stored representations (such as M_1, M_2 in Fig. 1a). More generally, S may be a source pattern in one cortical area, and M_1, M_2 target representations in another. The search for a linking sequence is bi-directional, starting simultaneously at the top (patterns M_1, M_2) and the bottom (pattern S). The nodes in this schematic figure represent patterns of activity (and should not be thought of as individual neurons), and the arrows indicate how patterns activate subsequent patterns in the se-

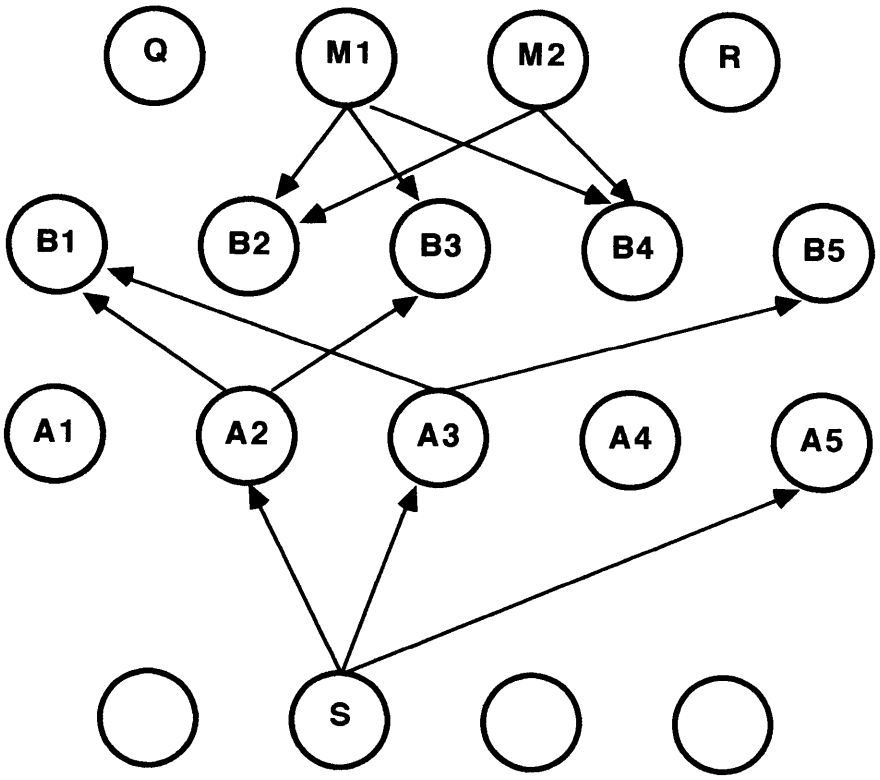

Figure 1a: The general sequence-seeking operation. A sequence of transformations linking the source pattern S with one of the target (Q, M1, M2, R) representations is sought. The search proceeds by expanding pattern sequences from both directions. Nodes in the figure represent patterns of activity; the arrows represent how patterns activate subsequent patterns in the sequence.

quence. For example, pattern S can activate A_2, A_3, and A_5, and pattern M_1 can activate B_2, B_3, B_4. In addition (and not represented in the figure), there are limitations on the co-activation of patterns; for example, patterns (B_2, B_3, B_4) may be prohibited from being co-active. In expanding the sequence down from M_1, only a subset of these patterns will be activated initially, and will later decay and be replaced by other members of this related group of patterns.

A linking sequence is successfully established when the two searches meet somewhere in this network of interconnected patterns. In Fig. 1a, a linking sequence can be established between S and M_1 via A_2, B_3. How can a successful link be found by the system and used, e.g., to select M_1 (that has been linked to S) over M_2 (that has not)?

The proposed scheme for linking the ascending and descending streams has two main components. First, the ascending and descending streams proceed along separate, complementary pathways. Second, when a track is being traversed in one stream, it leaves behind a primed trace in the complementary stream. This scheme is shown schematically in Fig. 1b. The figure is similar to Fig. 1a, except that each node is now split into two complementary nodes. For example, B_2 in Fig. 1a is now split into B_2 on the ascending pathway, and its complementary pattern $\bar{B}_2$ on the descending one.

The "priming" mentioned above has the following meaning in this structure. Suppose, for example, that the path from $\bar{M}_1$ to $\bar{A}_2$ along the descending path has been traversed. This will leave the complementary path $A_2 \rightarrow M_1$ along the ascending path in a primed state, ready to be activated. If an ascending sequence starting at S will later reach A_2, it will immediately follow along the primed path (to M_1), rather than explore alternative routes.

The bi-directional search now proceeds as follows. A number of sequences originating at S begin to be activated along the ascending pathway. At the same time, sequences originating at M_1 and M_2 begin to expand downwards along the descending pathway. Not all of the possible sequences are expanded simultaneously, but whenever a track (sub-sequence) is being traversed, the complementary track remains in a primed state. Suppose that by the time S has activated A_2 along the ascending stream, the track $\bar{M}_1 \rightarrow \bar{B}_3 \rightarrow \bar{A}_2$ had already been traversed in the descending stream. Due

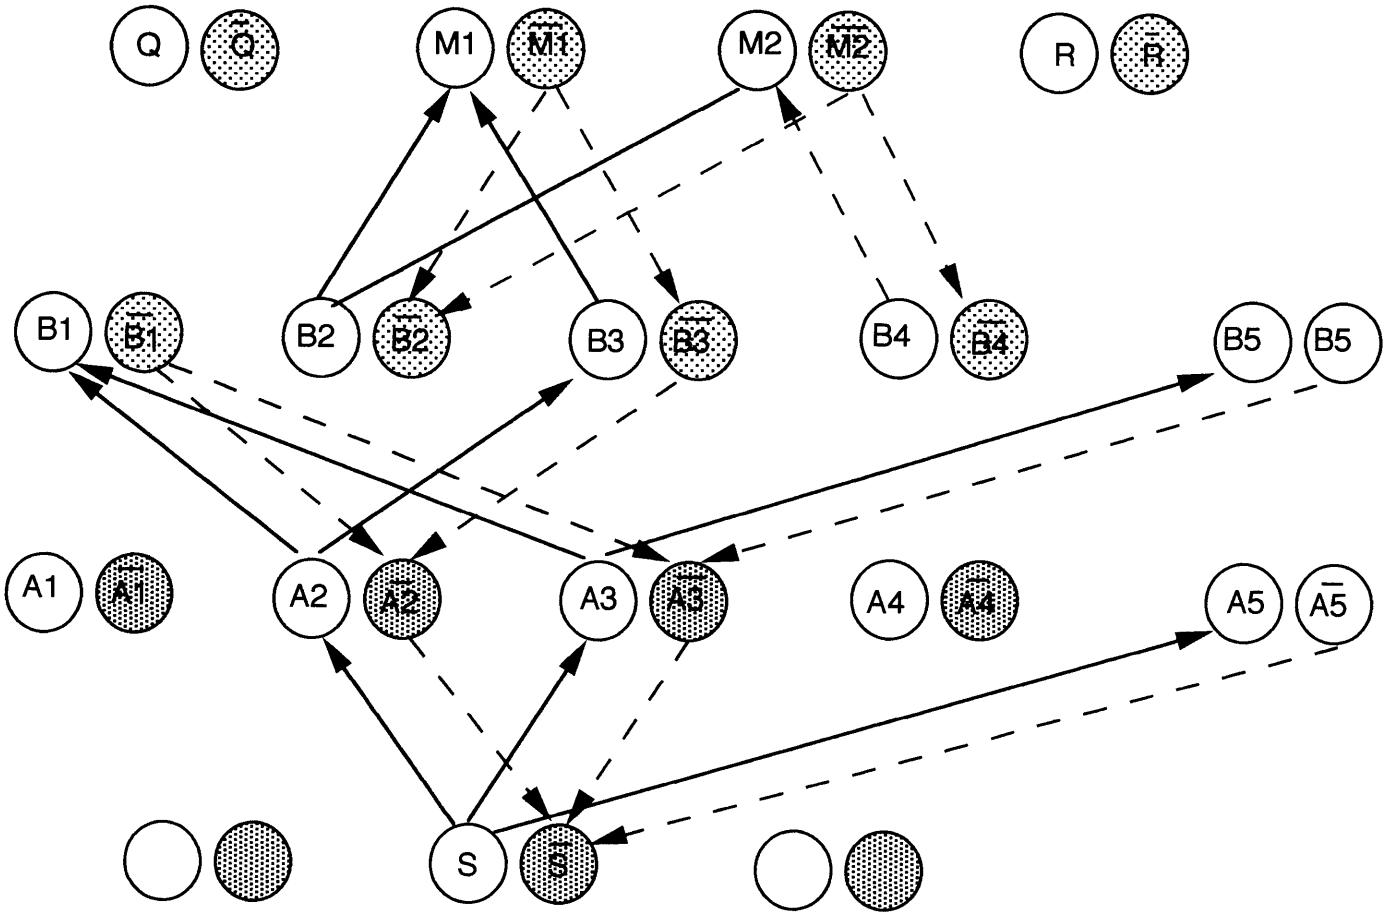

Figure 1b. The counter-stream structure. The figure is similar to 1a, but each node is now split into two complementary nodes. Unfilled circles and full arrows represent the ascending stream, shaded nodes and dashed arrows the descending stream. Every pair of complementary patterns is mutually connected by connections of the priming type (not shown).

to the primed traces, the result will be the immediate activation of the complete sequences $S \rightarrow M_1$ and $\bar{M}_1 \rightarrow \bar{S}$. In this manner, a complete link has been established between the representations (M_1 and S).

Two properties of this linking process are worth noting. First, a link between the ascending and descending streams can take place in this scheme at any level. The establishment of a successful link is accomplished by a combination of "bottom up" and "top down" processing, and the relative contribution of the two streams is flexible and can change with the circumstances. Second, if the meeting between the streams occurs at a particular node (A_2 in the example above), the ascending and descending patterns need not arrive at this node simultaneously. A meeting is also possible between an active pattern and a pattern that had been active some time before and decayed, but left a primed trace in the complementary stream.

Fig. 2 summarizes schematically the basic link in this chain of bi-directional processing. As before, nodes in this scheme denote patterns (i.e., activated sub-populations of units). On the ascending pathway, pattern A activates pattern B, and along the descending path pattern $\bar{B}$ activates $\bar{A}$. Corresponding patterns (A and $\bar{A}$, B and $\bar{B}$) are reciprocally connected by connections of the priming type. This basic link is embedded in a richly interconnected network in which each node is connected not to a single, but to multiple target patterns.

In summary, the basic design of the sequence-seeking model is relatively straightforward, comprising of two essentially feed-forward complementary networks going in opposite directions, with interaction between them primarily (but not exclusively) in the form of enhancing patterns across the two streams.

The biological embodiment of this scheme is discussed in more detail in Part II. Essentially, the processing takes place in two complementary pathways, with excitatory connections along the pathways and "priming" connections across. The excitatory connections between patterns are reciprocal, obeying the following general rule: whenever A is connected to B, there is a back-connection from $\bar{B}$ to $\bar{A}$, with cross-connections (of the priming type) between A and $\bar{A}$ and between B and $\bar{B}$. The reciprocity of the connections is an inherent aspect of the model, and it is also a distinguishing feature of cortical connectivity.

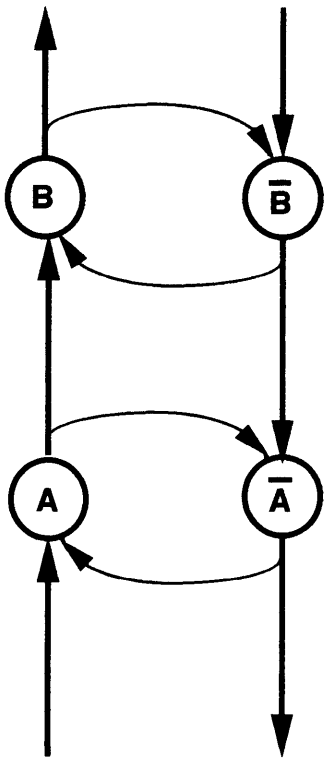

Figure 2. The basic unit of the counter-stream structure. Patterns A and B are on the ascending path, $\bar{A}$ and $\bar{B}$ on the descending path. The thin arrows between complementary patterns denote connections of the priming type. This unit is embedded in a network of richly interconnected patterns.

Learning sequences

The linking of the ascending and descending streams can be used to learn successful sequences. Since every pattern node in a successful sequence will receive both a direct activation and a priming signal from the complementary track, the approximate coincidence of the two signals can be used to preferentially strengthen the successful sequence.

A simple learning rule that will reinforce successful sequences, therefore, is the following. If the direct activation (e.g., $A \rightarrow B$ in Fig. 2) and the priming signal from the complementary track (e.g., $\bar{B} \rightarrow B$ in Fig. 2) are temporally close, then the direct connection ($A \rightarrow B$) is reinforced. This rule is local (it only depends on the activation at B), and yet it is sufficient to reinforce preferentially only successful sequences (that form a complete uninterrupted link between the source and target patterns).

What will be the effect of this reinforcement procedure? As discussed above, a given pattern in the network can activate a number of alternative successors, not all of which can be activated simultaneously. Increasing the effectiveness of a particular link, such as $A \rightarrow B$, will cause B to be expressed preferentially when A is activated, before some of the competing alternatives. As a result, out of the huge number of potential alternative paths, sequences that have proven useful in the past will be explored with higher priority.

The number of alternative sequences that a realistic size network can explore simultaneously is large, but it is still only a fraction of the total number of possible sequences that can be represented in the network. It is therefore useful to increase through learning the effectiveness of sequences that have proven useful in the past. Because the number of sequences that can be explored simultaneously is large, there is not much risk of saturating the system, and it will make sense to reinforce sequences that have proven useful in the past, and explore them again in the future under similar circumstances.

The entire process then relies on the breadth of the search, combined with the preferences induced by past experience. Rather than using sophisticated procedures to analyze the differences between the patterns being compared, the process generates in parallel a large number of alternatives, preferring ones that proved useful in the past, and then selects the most

promising ones.

In the process of reinforcing successful sequences, changes due to learning are distributed throughout the system, and are not confined to high level centers specializing in learning (Sejnowski 1986). Recent studies of learning certain perceptual skills suggest that low level visual areas, including primary visual cortex, are indeed involved in the modifications that take place during the learning process (Karni & Sagi, 1991).

In addition to the learning of complete sequences, as above, the system may also be engaged in the learning of the individual mappings, that is, the basic steps comprising the sequences. This aspect of the learning is, however, outside the scope of the current discussion, since the focus is not on the specifics of individual processes, but on their overall common structure.

The effect of context

Context can have a powerful influence on the processing of visual information (as well as in other perceptual and cognitive domains). A pair of similar, elongated blobs in the image may be ambiguous, but in the appropriate context, e.g., under the bed, they may be immediately recognized as a pair of slippers.

Context effects can operate in the framework of the sequence-seeking scheme by a prior priming of some of the nodes. The effect will be similar to the mutual priming of the ascending and descending streams but with longer time scales. (Priming between the streams may last for tens to hundreds of msecs., context effect should last for considerably longer, up to minutes or hours.) Sequences passing through the primed nodes will then become facilitated. In the above example, the location of the blobs, under the bed, will prime patterns representing objects that are commonly found in that location, making slippers a likely interpretation.

The general notion of priming internal representations is a common one, but its effects in the framework of the sequence-seeking scheme are particularly broad. When certain nodes are activated, e.g., by noticing and identifying the bed in the image, they will initiate sequences of their own, and an entire set of patterns may end up in a primed state. Later on, a large number of possible sequences passing through a primed node will be facilitated, compared with the non-primed sequences. It is not limited to

directly increasing or decreasing the likelihood of a single match, but it can have indirect, wide-spread effects, by facilitating otherwise less favorable sequences. A context pattern A may help to bring about the activation of B, not because A and B have been directly associated in the past, but because A may have a sequence leading to some intermediate pattern C, and, later on, a currently activated pattern may have another sequence leading to B via the primed pattern C. (Context may possibly involve also inhibitory effects, making some of the paths less likely, but this will not be discussed here further.)

These general characteristics capture some of the fundamental aspects of context effects in humans. Humans' perception and cognition appear to have an almost uncanny capacity (which is extremely difficult to reproduce in artificial systems) for bringing in relevant context information in a broad and flexible manner. The broad, indirect effects may be reproduced in part by context effects in the sequence-seeking scheme.

Express lines

In expanding the descending sequences, how can the initial selection of models be performed? Clearly, to cut down the number of competing sequences in the descending stream, it would be useful to expand with higher priority alternatives that appear more promising. For example, in attempting to recognize an object, some models (a face, say) may become more likely than others on the basis of partial analysis, although it may not be possible to identify the individual. It would be advantageous under these circumstances to expand many face-related sequences, possibly at the expense of others. Such an effect can be obtained by "express lines," directly connecting low-level to higher-level nodes along the ascending stream. The express lines will activate (directly or indirectly) patterns on the descending path. This will initiate an expansion of sequences from the selected patterns. This selection of higher-level patterns can be viewed as invoking a hypothesis suggested by the data, but which has yet to be confirmed. A link to the ascending stream will still be required to confirm the hypothesis. Note that, unlike the priming interaction between streams, in the case of express lines the ascending stream can directly activate the descending one. Express lines could also use inhibition rather than facilitation; if the partially expanded sequences in the ascending stream render some higher-level

nodes unlikely, inhibitory "express lines" could be used to suppress their expansion in the descending stream.

The express lines provide one mechanism for "indexing" into the large number of models stored in memory, that is, for the initial selection of likely models on the descending stream. This initial selection of models is not only to a single "topmost" level, models at different levels along the descending stream can also be "indexed" and serve as the starting point for descending sub-sequences.

Another mechanism for model selection is provided by the effects of expectation and context. Knowledge about the current situation can lead to the activation or priming of a subset of models that will then become preferential sources for descending sequences. The set of active models will then be modified and refined throughout the sequence-seeking process, as described in the next section.

Refining the expansion

Favorable sequences could also be expanded preferentially in another manner, based on the degree of match between the ascending and descending streams. For this process of refining promising sequences the linking between the streams should act as a graded event, depending on the degree of match between corresponding patterns. That is, the assumption is that if the two streams lead to sufficiently similar patterns, then the combined sequence will be activated, and the activation will reflect the degree of match between the corresponding patterns. As a result, sequences with higher degree of match will have a higher probability of being expanded in subsequent steps.

This will lead to a process by which coarse matches are established first, followed by successive refinements. Initially, the space of possible sequences will be sampled coarsely, subject to the limitations on the number of co-active sequences. The two populations of partial sequences, along the two streams, will then interact, by mutually comparing patterns and reinforcing favorable ones. The more promising sequences (the ones leading to better matches between the ascending and descending sequences) will then serve as preferred starting points for exploring additional sub-sequences. In this manner, the sequence seeking process will seek modifications and refinements of the more promising sequences.

This process has some features in common with a family of optimization and search procedures known as "genetic algorithms" (Holland 1975), that have been shown in some recent evaluations to behave quite efficiently compared with alternative methods, including a number of relaxation processes, gradient descent, simulated annealing, and heuristic search (Brady 1985, Peterson 1990).

Constrained transformations

As the ascending and descending streams are being expanded, they may sometimes reach a partial rather than a complete match. For example, in an image of scissors, a complete match may be obtained for one blade but not for the other. More generally, if an object has been distorted in a flexible, non-rigid manner, a sequence of transformations may lead to a good match of a part of the object but not all of it. When this occurs, it would be advantageous to maintain the partially successful match in subsequent transformations of the object. The unmatched part will transform, but the matched part will remain "frozen." The question of how such constrained transformations may be obtained is an interesting one, but will not be considered further here. Let me note, however, that such a process could play a useful role in the recognition of objects undergoing complex, non-rigid distortions, a task that proved to be one of the most difficult aspects of theories of visual object recognition.

Controlling the sequence-seeking process

The sequence-seeking process is likely to include control mechanisms that help to restrict and coordinate the sequences being explored. Perhaps the most straightforward example of such a mechanism in visual perception is spatial selection, i.e., restricting the region in the incoming image that is being explored. There is substantial evidence in visual perception for the existence of such processes. A popular metaphor has been the "spotlight" of attention (Posner et al. 1980). Regions inside this "spotlight" are selected for more extensive processing. I will not review here the evidence and theories related to these selection processes. For the current discussion, the main point is that selection processes of this type will serve to select and control the expanding sequences in the ascending stream. In a biological implementation, discussed in Part II, it is suggested that a special control

loop, possibly involving layer 5 of the cortex and subcortical structures may be involved in this control process.

Different types of relations

In the sequence-seeking scheme, a sequence is expanded when a pattern currently at the tail end of the sequence activates new patterns connected to or associated with it. Connections between patterns may be used to signify that the two are related, but relations come in different types. For example, the notion of a "dog" is associated with specific kinds of dogs, but also with entirely different notions (such as a leash). When attempting to recognize a dog, some of the connections are relevant and should be activated, while others are not. This is a general problem in many models; it is not enough to associate different representations, it is also important to make at least some distinctions between different types of associations. What can distinguish in a network between different types of association?

One common solution is to assume that connections somehow come in different "flavors" or types. There is one type specifying, e.g., being-part-of, another signifying sub-categories, etc. Different types of connections can be readily implemented in artificial systems, but are more difficult to envisage in a biological setting.

I will not discuss possible solutions in detail, partly because I suspect that there may not be a single simple solution. Instead, I will mention briefly two possible aspects of the solution. Part of the solution may be that different types of relations are implemented in different sub-networks in the system. For example, there may be one network in one cortical area where the connections can be interpreted as specifying sub-categories, whereas "being-part-of" or "sounds-like," etc. are implemented in different sub-networks and different cortical areas. One of the tasks of the control processes mentioned in the previous section may then be to enable and disable sub-networks as required.

A second aspect of the solution may come from the use of spatial s-election, discussed in the previous section, as illustrated by the following example. It was mentioned above how context information may affect the flow of the expanding sequences, helping to identify, for instance, a pair of elongated blobs under the bed as a pair of slippers. In this example, "bed" and "slippers" are associated in memory, but it is a special type

of association. The notion of "bed" is also associated with a large number of other notions such as "sofa" and "furniture" which, unlike slippers, are not associated with being under the bed. One can imagine a useful scheme in which a pattern representing a bed is associated as a whole in an associative memory with a number of other patterns related to "bed." In addition, a sub-pattern corresponding to the location under the bed is associated, e.g., with the notion of slippers. Here is where spatial selection can play a useful role. If during the processing the entire "bed" pattern is activated, various patterns associated with "bed" will become activated next. If, however, the sub-pattern representing the region under the bed is selected, items associated with being under the bed will be selected next. As this example illustrates, connections between sub-patterns can be used to specify different types of connections.

In concluding this part, it will be useful to follow in general outline the flow of the sequence-seeking process, for example in the case of the visual identification of a familiar shape. The processing and analysis of the input pattern will center in this scheme around the bi-directional search for a linking sequence of mappings between a source and a target representation, and it will include the following stages. The input pattern will initiate sequences of transformations and mappings, expanding in a bottom-up manner. Rather than a single stream of processing stages, multiple alternatives will be explored simultaneously. These sequences will include geometric transformations, such as changes in position, orientation, and size, and also transformations and mappings that are not geometric in nature, such as exploring alternative ways of representing and abstracting the input shape. The expanding sequences will remain active for a while, then decay and be replaced by alternative ones. The sequences being explored will leave the complementary descending tracks in a primed state.

At the same time, sequences will be expanded in a top-down manner, starting from previously stored patterns. The transformations applied to the stored representations will be shaped by past experience with these patterns. The descending sequences will also leave their complementary tracks in a primed state.

Some of the descending sequences will be favored by 'express lines' from the ascending patterns, and will be expanded preferentially, based on some

properties of the shape, such as its color, overall elongation, curvature properties etc. The expansion will also be biased by context effects (e.g. by the priming of sub-sequences), and by the rough agreement between ascending and descending patterns. If a partial match is obtained (i.e., a close match between subpatterns), it will constrain further transformations, restricting them to the non-matching parts. The control loop will influence the process by selecting the patterns to be expanded and the sub-structures participating in the computation.

A close match between corresponding patterns will lead to the establishment of a complete link between the incoming and a stored representation. The connections participating in the link will be reinforced to facilitate the future processing of similar patterns.

PART II
Biological Implications

This part will focus on the biological implications of the sequence-seeking scheme and its implementation in the counter-stream structure. The discussion is divided into two parts. The first examines the implications of the model for the general connectivity patterns of the cortex, both between different areas and between layers of the same area. The second discusses properties of the different types of connections in the model. The model will be compared with known connectivity patterns in the visual cortex, where a large number of connections have been studied in detail. There is evidence, however, that other cortical areas, including somatosensory, motor, and auditory, may exhibit generally similar connectivity patterns (Felleman & Van Essen, 1991).

2.1 Connectivity Patterns

The connections between cortical areas, as well as between layers within an area, will be examined and interpreted in this section in light of the sequence-seeking model discussed in Part I. As we shall see, the proposed scheme appears to be in general agreement with the existing data, and it also gives rise to a number of new predictions. The discussion involves some fine details of the anatomy, but it serves to illustrate the type of predictions,

on both general and specific levels, that can be derived from the model.

The basic connectivity scheme

A skeleton of the basic connections in the sequence-seeking scheme is shown in Fig. 3a. It is composed of two counter streams, an ascending one ($A \rightarrow B \rightarrow C \rightarrow D$) and a descending one ($\bar{D} \rightarrow \bar{C} \rightarrow \bar{B} \rightarrow \bar{A}$), with cross connections (of the priming type) between complementary structures. In a biological implementation, each of the patterns (A, B, etc.) is assumed to be a sub-population of neurons, acting together, possibly with some degree of synchrony. A connection $A \rightarrow B$ means, therefore, that a certain sub-population of neurons representing pattern A, when activated together, will activate another sub-population B, usually in a different layer or a different cortical area. The next section discusses aspects of cortical connectivity in light of this general structure. It should be noted, however, that the diagram illustrates schematically the activation of patterns and the relation between complementary patterns. The actual implementation of this scheme in terms of populations of neurons is likely to be considerably more complicated.

In terms of cortical interconnections, the model suggests the schematic connectivity pattern shown in Fig. 3b. There are again two interconnected counter streams. The ascending one goes through layer 4 to a sub-population of the superficial layers, denoted in the figure as AS (for Ascending Superficial). It then projects to layer 4 of the next cortical area (II in the figure). The descending stream goes through a different sub-population of the superficial layers (DS, for Descending Superficial) to layer 6, and from there to DS of a preceding area. The connections can also leap over one step in the stream, e.g., AS directly to AS on the ascending stream, and DS $\rightarrow$ DS or 6 $\rightarrow$ 6 on the descending stream (thin lines in Fig. 3b). Layer 5 is not depicted in the diagram because, according to the model, it is involved primarily not in the main streams, but with their control, via sub-cortical areas such as the pulvinar.

There are two reasons for the assumption that layer 5 (or parts of it) may be involved in the 'control loop' mentioned in Part I. The first reason is that layer 5 (or parts of it, such as sublayer 5B in the macaque's primary visual cortex) projects in an orderly topographic manner to a number of subcortical structures, such as (in the visual cortex) the pulvinar, and the

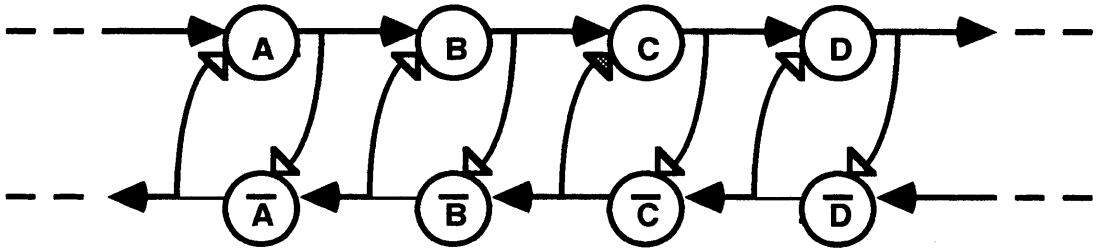

Figure 3a. The main connections along and between streams. Full arrows denote the connections along the stream, unfilled arrows the priming connections between corresponding patterns in the complementary streams.

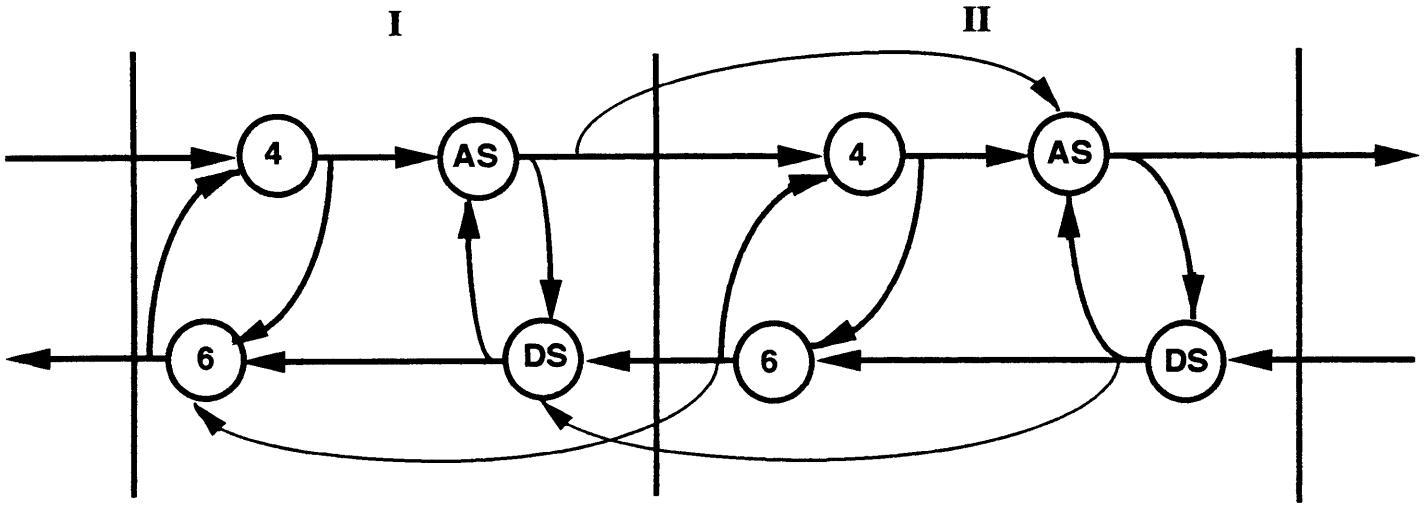

Figure 3b. The basic counter-stream scheme, showing the main connections within and between cortical areas. The ascending path goes through layer 4 and the ascending superficial population (AS) to the next area. The descending path goes from the descending superficial population (DS) to layer 6 and to the lower area. The connections between the streams are of the priming type. The thin arrows show pathways that "leap over" one step in the stream.

superior colliculus. The pulvinar is reciprocally connected in a topographic manner to multiple visual areas. This pattern of connectivity suggests that some of the subcortical areas may be involved in the integration and control of the activity in different visual areas, and that layer 5 is a part of this control loop.

The second piece of evidence that might implicate layer 5 with some control functions is the finding that many layer 5 pyramidal neurons show a special characteristic firing pattern (Silva *et al.* 1991, Kasper *et al.* 1991). These cells were shown to have a tendency to produce rhythmic and locally coordinated firing pattern. Silva *et al.* suggested on the basis of their observations that "...neurons in layer 5 can initiate synchronized rhythms and project them on neurons in all layers" (p. 434). The anatomical and physiological evidence is only partial, but it suggests that layer 5 cells may play a control and coordination role.

The scheme diagrammed in Fig. 3b implicates layer 6 with the ascending stream and layer 5 primarily with a separate control function. This division may in reality be less clear-cut. For example, parts of layer 5 may be situated on the descending path (Ferster and Lindström 1985), and parts of layer 6 may perhaps participate in the control loop. (For example, a population of layer 6 cells projects in a topographical manner to the claustrum, which is in turn reciprocally connected to multiple visual areas.) It may be more appropriate, therefore, to refer to the role of layer 6 in the diagram as "DI", for "Descending Infra", similar to the DS and AS populations in the superficial layers. The diagram is also schematic and over-simplified in other respects. For example, there may be additional steps within an area (as in the parvo stream in V1). However, the main goal of the diagram is not to account for possible variations, but to emphasize the common underlying structure according to the model.

Connections of V1: data and predictions

Figure 4a,b shows a more detailed version of this diagram, applied to cortical area V1 (for which the data are more complete than for other visual areas), and its connections to the LGN and cortical area V2. Fig. 4a shows the connections in the macaque of the magno-cellular stream, and 4b of the parvo-cellular stream, on the way from the LGN (where the ascending

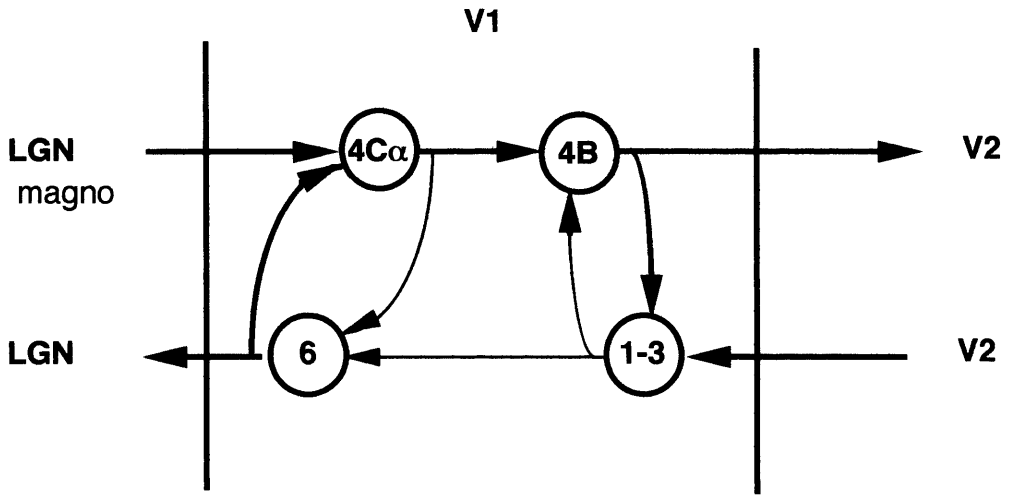

Figure 4a. The main connections according to the model along the two streams from the LGN (magno) through area V1 to V2. Thick arrows: established connections. Thin arrows: connections predicted by the model.

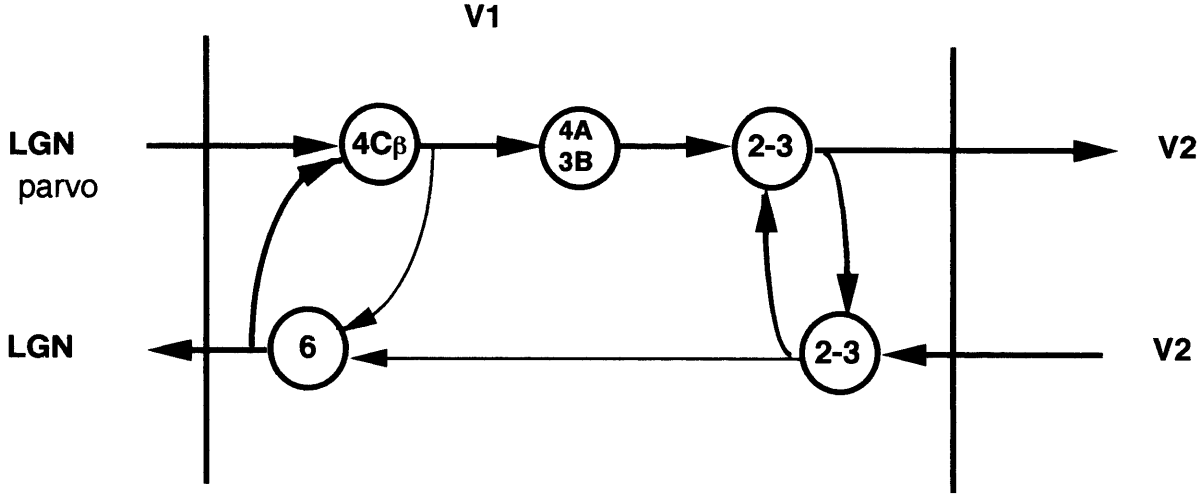

Figure 4b. The main connections according to the model along the two streams from the LGN (parvo) through area V1 to V2. Thick arrows: established connections. Thin arrows: connections predicted by the model.

and descending streams in fact converge) to area V2 (e.g., Rockland and Lund 1983, Lund 1988a, 1988b, Martin 1988). The division into the two streams may be somewhat schematic; there is evidence (Lund 1988a) for some degree of mixing between them.

The diagram shows the main connections; there is evidence for additional secondary ones that will not be discussed in detail. They include: (i) sparse magnocellular projections to the lower half of layer 6 (Blasdel & Lund 1983), (ii) a parvocellular projection to 4A with collaterals also to 3B (Fitzpatrick *et al.* 1983, Blasdel & Lund 1983), (iii) some parvocellular input to the upper part of layer 6 (Hendrickson *et al.* 1978), with possible collateral to layer 1 (Blasdel & Lund 1983), (iv) projection to layers 2-3 of the cytochrome oxidase-rich blobs, (v) some possible projection from the upper part of layer 6 and lower 5 to V2 (Kennedy & Bullier 1985).

The diagram includes both known connections and connections predicted by the proposed scheme that have not been described in the literature. The predicted connections are discussed in the next three subsections. The first relates to a connection within the descending stream, the next two to the reciprocal connections between the two streams.

The Connections from the Superficial Layer to Layer 6

The model predicts an excitatory connection from the superficial layers to layer 6. This is a part of the descending stream in the model, from population DS to layer 6 in Fig. 3b. Such a connection appears physiologically possible, since the apical dendrites of layer 6 pyramidal neurons projecting back to the LGN extend to the superficial layers (Martin 1988, Fitzpatrick and Einstein 1989) where they can make connections with the DS population.

Interconnections in the Superficial Layer

In the magnocellular stream, the model predicts reciprocal interconnections between layer 4B (playing the part of AS in the model), and layers 1-3, the recipients of descending projections from V2 (DS in the model). The projection from 4B to the superficial layers is well established. It is also known (Lund 1988a) that 4B pyramidal cells send apical dendrites to the superficial layers where the connection may take place.

Lund (1988a) notes a puzzle regarding the connections between AS and the superficial layers: cells in 4B are strongly directionally selective, but the superficial layers that receive input from them are markedly less directionally selective. The superficial layers show strong color coding, but cells in 4B do not. Such discrepancies are in fact compatible with the counter-streams scheme. According to the model, these two populations are situated on complementary streams, with priming connections between them. It is therefore not expected in this scheme that the prevalent properties of one of these populations will be strongly determined by the other. The superficial layers also contain other sub-populations, in particular, the AS population of the parvocellular stream. In the model, the color-coded non directional cells in the superficial layers are likely to be a part of this stream.

There is another specific prediction suggested by the model regarding these connections. As can be seen from the diagram, the expectation from the model is that the superficial cells connected to 4B will also be the recipients of descending projections from V2.

Interconnections Between Layers 4 and 6

The counter stream model in Fig. 4a, 4b, predicts a reciprocal connection between layer 4, and the LGN-projecting cells in layer 6.

Layer 4C in the macaque is known to receive excitatory connections from layer 6 (Lund 1988a); similar connections are also known to exist in the cat primary visual cortex (McGuire et al, 1984). There is also evidence for the existence of connection in the opposite directions, as reviewed briefly below.

In the tree shrew, connections of this type (from 4A/3B) have been reported, (Lund *et al.* 1985). In the magnocellular stream of the macaque, Lund & Boothe (1975) found that there is a population of layer 6 pyramidal cells with arbors of their apical dendrites in $4C\alpha$ where they may collect input from axons of layer $4C\alpha$ cells, and also layer 6 pyramidal cells with dendritic arbors in 4B, the next station in Fig. 4a along the ascending stream. They also report axonal projections from layer 4B to layer 6. This strongly suggests that there are indeed anatomical routes for connecting two stations along the ascending stream ($4C\alpha$ and 4B) to layer 6.

In the parvocellular system, they report layer 6 pyramidal cells with dendritic arbors in 4A and 3B, sending axons back to these layers, with collaterals in $4C\beta$.

This manner of reciprocal connections between layers 4 and 6 is consistent with the requirements of the model. The main role assigned to these connections by the model is to provide cross connections, presumably of the priming type, between the two streams.

It also interesting to note in this regard that the population of layer 6 cells projecting back to the LGN, were found (in the cat) to be the same cells that are also connected to layer $4C$, by axonal collaterals and dendritic arbors (Katz, Burkhalter, & Dreyer 1984), in accordance with the model in Figs. 4a,b.

General connectivity pattern: data and predictions

The previous section focused on the connectivity of the primary visual cortex. This section examines the general connectivity pattern within and in between areas in light of the sequence-seeking scheme.

Connections Within an Area

If the hypothesis regarding the counter-streams structure is broadly correct, then a number of predictions can be made regarding the main connectivity patterns within an area. The predictions listed below are for the most direct implementation of the counter-streams structure, some modifications and exceptions are possible in some cases.

One prediction is the distinction within the superficial layers between the AS and DS sub-populations. The AS population on the ascending stream projects to a higher visual area, and the DS population on the descending stream is the target of back projections from the higher area. A number of different AS and DS populations will exist in parallel within the same area, e.g., projecting to (and receiving from) distinct targets. This scheme stands in contrast to an alternative view (Sandell & Schiller 1982, Maunsell & Van Essen 1983, Rockland & Virga 1989) in which the back projections contact and modulate their own source.

A separation between the two populations along the lines suggested above reflects the most straightforward implementation of the scheme. Such

a separation is not easy to establish empirically, but, as reviewed below, at least in some cases that can be examined, it seems to be honored by the system. It should be noted, however, that some exceptions can exist without violating the constraints of the model. For example, priming connections can connect a population on one stream with a population on the complementary stream in a different visual area. As a result, an AS population, for instance, can be both the source of ascending connections, as well as the recipient of (priming) connections from the higher area. If such connections exist, they will make the empirical testing of the separation more difficult.

The separation between ascending and descending populations is evident in the connections involving layer 4: the ascending streams terminate in layer 4, the descending streams always avoid it. In the superficial layers the situation is more difficult to assess, and the available evidence is at present restricted.

An example involving the superficial layers is the magnocellular projection from V1 to V2. Rockland and Virga (1989) find it surprising that there are no terminations of back projections from area V2 to layer 4B of V1 (the source of efferents to V2). As illustrated in Fig. 3b, this is in fact to be expected in the current model. It is further expected that even in those cases where the same layer is both the source and the target of connections to another area (such as the connections between layer 4B of V1 and MT), there are likely to be two distinct sub-populations, the AS and DS, within the layer. The projection from layer 4B of V1 to MT was studied by Zeki and Shipp (1988). Retrograde labeling revealed a patchy origin, suggesting that only a sub-population of 4B neurons participate in the projection to MT. It would be of interest to know whether the back projection from MT is directed primarily towards the complementary sub-population.

A second prediction concerns reciprocal connections within an area. That is, the AS and DS subpopulations within the same area are expected to be connected in a reciprocal manner. Similarly, layer 6 (or a sub-population on the descending path) is expected to be connected reciprocally with the ascending path, in particular, with layer 4. These connections are expected to be of the priming type (Section 2.2).

A third prediction is that a connection is expected on the descending path between the superficial layers (sub-population DS) and layer 6 (or

the DI population). (This connection may become less important when there is a direct connection from layer 6 of the higher area to layer 6 of the lower area, one of the "leaping" connections in Fig. 3b). This connection is expected to be of the "standard" excitatory type (Section 2.2).

Connections Between Areas

We next examine the general patterns of connectivity between cortical areas. It will be convenient in examining these connections to follow the classification of projections between areas into forward, backward, and lateral connections (Rockland & Pandya 1979, Maunsell & Van Essen 1983, Friedman 1983, Van Essen 1985, Andersen *et al.* 1990, Boussaud *et al.* 1990, Felleman & Van Essen 1991). It has been observed in these and other studies that the connections between cortical areas (not only visual, but also somatosensory and motor) can be classified into "forward," "backward," and "lateral" connections, on the basis of the laminar distribution of their source and destination. Roughly speaking, the distinction is the following (more details are given below). Connections of the forward type terminate in layer 4, and (with a few exceptions) originate in the superficial layers. Backward connections avoid layer 4, and terminate instead in the supra and infra layers. Lateral connections terminate in all layers (layer 4, as well as supra and infra).

When the connection from a given area A to a second area B is of the forward type, the connection from B to A is almost invariably of the backward type, and area B can then be considered 'above' A. This ordinal relation can serve to suggest a hierarchy of cortical areas, which has been recently reviewed in detail by Felleman and Van Essen (1991). In the sequence-seeking scheme, there is also a distinction between a "forward" connection (along the ascending stream), and a "backward" connection (along the descending one). It is also possible in the scheme for these connections to co-exist (i.e., areas A and B can be connected by both forward and backward connections), which will give rise to lateral connections.

Although in visual processing there is a natural general hierarchy of processing from the incoming signal to higher, more abstract levels, the sequence-seeking model by itself does not require strict hierarchy, and can exist in a more general structure of reciprocal connections. Violations of

the strict hierarchy are possible, and are more likely to occur at the higher, more abstract levels of processing. In the discussion below, the notion of a hierarchy does not play a crucial role, but the classification of the connection types into forward, backward, and lateral, serves a useful role in organizing the discussion. As we shall see, the somewhat complex set of possible projection patterns can by and large be deduced from the counter-streams structure diagrammed in Fig. 3b. In addition, this structure may help to resolve a number of apparent exceptions to some general connectivity rules.

The Forward Connection

The typical forward connection between cortical areas is from the superficial layers of one area to layer 4 of the next area (Van Essen 1985). This pattern of connections is consistent with the ascending pattern in Fig. 3b. The figure also shows the possibility of a connection that leaps over one stage in the stream, from AS of a given area to AS of another (in addition to the projection to layer 4). Such connections have indeed been observed (e.g., Rockland 1990; the parvocellular projection to layer 4A of V1 may also be of this type).

In addition to this main pattern, some forward connections have been found to originate from both the supra and infra layers (Andersen *et al.* 1990, Felleman & Van Essen 1991). The ascending projection from the infra layers is not a part of the counter-streams structure of Fig. 3b. There are several possibilities to account for such a connection within the counter-streams pattern of connections. For example, it may be that the connection is in fact a lateral rather than a forward one (see also the discussion of the lateral connection below). A lateral connection will contain an ascending connection as well as a descending one, that will supply the infra $\rightarrow$ superficial component. The result will be a bi-laminar connection (from both supra and infra layers) of one area to layer 4 and the DS layer of another. This may be the case, for example, in the projection from V3A $\rightarrow$ MT, which has the unusual forward connection, and where the opposite projection MT $\rightarrow$ V3A is suspected, in fact, to be of the lateral type.

What should be excluded, according to the model, is that the back-projecting neurons of layer 6 will also be the origin of an excitatory ascending projection. (They could, however, be the source of priming connections

to the ascending stream.)

The Backward Connection

The backward connection is characterized by Felleman & Van Essen (1991) as having one of two patterns, which they denote as I → M or B → M. (I denotes infra, B denotes bi-laminar, and M denotes termination that avoids layer 4.) The first type of connection is from the infra layers (I) to both infra and supra layers, but avoiding 4 (M). The second type is from both infra and supra layers (B), again to M (avoiding 4). In the counter-streams pattern of connections (Fig. 3b), the back projection is composed of the connection 6 → DS, and the "leaping" connections (DS → DS, 6 → 6, and perhaps also DS → 6). This will give rise to both of the patterns listed by Felleman and Van Essen (1991). It also suggests possible additional sub-cases, for example a connection 6 → DS only. The back projection from V2 to V1 in the macaque may be of this type. Reconstructions by Rockland and Virga (1989) show projections to the superficial layers (mainly 1 and 2), with restricted collaterals in layer 5 (but not in 6).

The Lateral Connection

Lateral connections in the hierarchy detailed by Felleman and Van Essen (1991) are projections that do not fit the forward or backward types. In their terminology, the pattern of lateral connections is B → C; the origin is bi-laminar, both supra and infra layers (B), and the target is columnar, spanning all layers (C). The pattern of connections appears complex and sometimes perplexing, but many aspects of it can be derived almost directly from the general counter streams-scheme. In the counter-streams model lateral connections are in fact expected, and, in terms of the pattern of Fig. 3b, the suggestion is that the lateral connection is simply the "union" of ascending and descending connections. If this view is correct, then the main connections participating in the lateral connection can be inferred immediately from the basic scheme (Fig. 3b). They include the direct connections:

- AS → 4
- 6 → DS

and the connections that leap over one step:

- AS $\rightarrow$ AS
- DS $\rightarrow$ DS
- 6 $\rightarrow$ 6

It can be seen that the origin of the projections would indeed be bilaminar, and the terminations would span all layers. If this view of the lateral connection is correct, it gives rise to some predictions, and offers a possible explanation for difficulties noted in the past, and discussed below.

The Problem of Irregular Terminations

This difficulty arises when a projection from one area A to another area B shows a mixed pattern of terminations, that are incompatible, according to the hierarchical view. For example, when some of the terminations are restricted to layer 4 of area B, while other labelings show columnar terminations. This mixed pattern is denoted F/C by Felleman and Van Essen (F for termination in layer four, C for columnar termination). Termination in layer 4 only is a signature for connections of the ascending type, but a columnar termination signifies a connection of the lateral type (Felleman & Van Essen 1991). How, then, to classify the $A \rightarrow B$ connection in this case?

The view from the counter-streams structure suggests that the connections from the superficial layers of A to area B in a lateral connection are composed of two projections:

- ascending: AS $\rightarrow$ 4
- descending: DS $\rightarrow$ DS, 6

(In addition, there is a descending connection 6 $\rightarrow$ DS.)

Anterograde labeling of the upper layers of area A can therefore show a mixed pattern of terminations. (In anterograde labeling, dye is carried forward along cells axons and across the synapses to the post-synaptic cells.) The termination pattern may be either in 4 alone (from the AS population

of A), or a columnar termination (from AS and DS). This may be the F/C paradoxical termination noted by Felleman and Van Essen. It can also (by labeling the DS sub-population alone) show a bi-laminar pattern of connections in both the supra and the infra granular layers, but avoiding layer 4. This can account for the other types of irregular terminations.

A related prediction is that in a lateral connection between areas A and B, two distinct sub-populations should exist in the superficial layers of B. One is the AS, the origin of ascending connections to A, the second is the DS, the target of descending connections from A. The two should be distinct and interconnected (by connections of the priming type).

Origin-Termination Mismatch

An example of this difficulty (e.g. from area V4 to MT) is the following. In the connection from area A to B, the termination is columnar (all layers), but the origin is sometimes bi-laminar (both supra and infra), and sometimes only superficial (Felleman & Van Essen 1991). The problem is that columnar termination signifies a lateral type of connection, which is inconsistent (according to the hierarchical classification) with an exclusively superficial origin. The resolution suggested by the counter-streams configuration is that the connection is in fact of the lateral type. Retrograde labeling of the superficial layers of area B can label A either in a bi-laminar fashion (from the DS sub-population of B, which receives input from both layers, Fig. 3b), or just the superficial layers (from the AS sub-population, which receives input from superficial layers only, Fig. 3b).

Finally, it is worth noting that the counter-streams system can withstand the elimination of some connections without affecting significantly its sequence-seeking function. For instance, it is possible to have connection of the lateral type in the projection from area A to B, and at the same time the reciprocal connection from B to A could be of the descending (or ascending) type only. (One can expect a higher probability of removing the ascending link in the higher level areas and the descending one in the lower areas.)

Why six layers?

The counter-streams scheme suggests a rough functional account for the

separation of the cortex into five-to-six main layers (five, plus the somewhat special layer 1). Roughly speaking, within this model, two are required for the ascending stream, two for the descending stream, and one for the control loop. (Layer 1 may participate in the control loop as well as in making some of the intrinsic connections within the streams.) The layers will contain sub-layers, e.g. different sub-populations of AS, projecting to different targets on the ascending stream, and similarly for the DS population on the descending stream.

The division is not precise, because, for instance, (i) the two functional superficial populations (AS and DS in Fig. 3b) may not be always segregated in a consistent manner, (ii) sometimes (as in the ascending parvo stream in V1) extra sub-layers may be required, (iii) layers 1 and 5 may be involved in both the control function and in providing other connections. However, variations and elaborations are naturally likely to exist in such a complex system, and they may not be inconsistent with this basic account.

2.2 Properties of the Connections

The previous section examined the "wiring diagram" implied by the sequence-seeking model. In this section, I will examine properties of different connections implied by the model.

The Role of inhibition

The model makes specific predictions concerning the excitatory connections between patterns. However, there is also a clear role for inhibitory connections. One role of inhibition is within the basic mappings, or transformations, comprising the ascending and descending streams. Marr, in his models of the neocortex (Marr 1970) and archicortex (Marr 1971), suggested two functions for inhibitory connections in the mapping between two patterns, represented by two populations of neurons. One function is to ensure that the number of co-active output units is approximately constant. This is obtained by affecting their threshold, using a mechanism of divisive inhibition (implemented by basket cells). The second function is to adjust the dynamic range of the output, and this is obtained by subtractive inhibition. Such inhibitory connections could serve similar purposes in the counter-streams architecture.

In addition to their role within the individual mappings, inhibitory connections are likely to be required for three additional roles. The first is to prevent the co-activation of too many patterns (see also appendix). A second possible function is inhibition between incompatible sequences. For example, patterns derived in the ascending stream may be sufficient to rule out certain possible interpretations, and it would be useful to inhibit unlikely sequences from expanding further in the descending stream. Finally, inhibition may play a role in the control loop to suppress some of the sequences from expanding further.

Connection types

The counter-streams structure uses excitatory connections for a number of different functions. The excitatory connections include (i) lateral connections among neurons belonging to the same patterns, (ii) connections between successive patterns in a stream, and (iii) priming connections between complementary patterns in the opposing stream. This is consistent with the observation that the thalamic inputs to the primary thalamic recipient cells in the cortex account for only the minority, perhaps 25%, of the total excitatory inputs to a single neuron (Lund 1988a).

An important requirement in the model regarding these connection types is the existence of priming connections between the complementary streams. In this section, I discuss briefly the evidence and possible mechanisms for these connections.

The requirement from the priming connection is that whenever a pattern A in the ascending stream is activated, it will "prime" the corresponding pattern $\bar{A}$, making it easier to activate for a certain time period following the stimulation. Similarly for the descending stream, $\bar{A}$ will activate its counterpart A. Priming interactions of this predicted type have not been investigated directly, but a number of reported findings appear to be consistent with such a mechanism.

Ferster and Lindström (1985) recorded (both extracellularly and intracellularly) from layer 4 in the cat's primary visual cortex during an electrical stimulation of the LGN. With low-amplitude, low-frequency (1 Hz) stimulation, the response was mediated only by the geniculo-cortical projection (from the LGN to layer 4). At a higher frequency and amplitude,

the stimulation also activated layer 6 cells by antidromic activation. These cells, via their collaterals, add excitatory input to the primary layer 4 cells. These layer 4 cells now receive two kinds of inputs – a direct one from the LGN, and a second one from the collaterals of layer 6 cells. Under these conditions, the response is enhanced significantly, the probability of firing increases, and most cells fire multiple spikes in response to each stimulus. Under the opposite conditions, when layer 6 is inactivated, the main observed effect is the reduction in excitability of layer 4 cells (Grieve, Murphy & Sillito, 1991). These findings support the possibility that the connection from layer 6 to layer 4 is of the priming type in the model, and can enhance the response of layer 4 cells in the ascending stream.

Some indirect support for the possible priming of layer 6 cells by the ascending stream comes from the source-density analysis of Mitzdorf and Singer (1978) in visual areas V1 and V2 in the cat. This analysis indicated that the activation of the superficial layers by an ascending stimulation did not cause significant activation of layer 6, in accordance with the general separation of ascending and descending streams. The main observable effect of the ascending stream on layer 6 (in addition to some direct LGN input) appeared to be a long-lasting subthreshold excitatory input. This is consistent with the possibility of a priming interaction (see also the discussion of possible mechanisms below). There was also evidence for long-lasting subthreshold excitatory input in the superficial layers, which, in the model, is attributable to the priming effect of the ascending stream on the complementary DS population.

Interestingly, in a similar study in the monkey (Mitzdorf & Singer 1979), a long lasting subthreshold excitation was also apparent in layer 4B of the primary visual cortex. This is attributable in the model to the reciprocal connections between layer 4B in area MT. According to the model, layer 4B contains two sub-populations , one on the ascending stream to MT, the other on the complementary descending stream. Ascending stimulation will then cause a direct excitation of the ascending sub-population, followed by a long-lasting excitation of the complementary descending population.

In another study possibly related to the different excitation types, Albowitz and Kuhnt (1991), using voltage sensitive dyes, observed two types of excitation patterns in the visual cortex: a fast one, followed by hyperpolarization, and a slower, gradual one, without ensuing hyperpolarization.

These may correspond to the distinction between the direct and priming excitatory drives.

Finally, a finding regarding priming between the streams comes also from studies of the priming of LGN neurons by projections from the cortex. Sur and his collaborators (Esguerra & Sur 1990) have examined the response in LGN slices from the ferret to ascending and descending stimulations (from the optic tract and optic radiation, respectively). They found that when a descending stimulation preceded an ascending one, it enhanced the response to the ascending stimulation in a pronounced, non-linear manner. The intracellular responses (EPSP) were larger under these conditions than the algebraic sum of the individual responses, longer in time, and with higher probability of driving the cells. This is, again, consistent with the priming interaction in the model between the descending and ascending streams.

Possible mechanisms

There are reasons to believe that the priming of one projection by another is a plausible neuronal mechanism. Evidence for two possible mechanisms of this type has been reported. One mechanism has been investigated by Miller, Chapman, and Stryker (1989). In this study, responses of cells in the cat's visual cortex to visual stimulation were profoundly suppressed by the blocking of NMDA receptors (by using APV). A possible mechanism proposed by Miller *et al.* by which NMDA receptors could control the responsiveness of cells is that such receptors, when activated in neocortex pyramidal cells, cause a slow, long-lasting EPSP that rises to a peak in 10-75 msecs. As a consequence, "such a slow EPSP could provide a base upon which continuing subthreshold input, mediated by non-NMDA receptors, would become suprathreshold" (p. 5187).

Another mechanism has been proposed by Koch (1987) and has received support from studies by Sur and his collaborators (e.g., Esguerra, Kwon, & Sur 1989) and by Sherman *et al.* (1990) in the LGN. An important property of NMDA receptors is that they can increase the cell's response in a nonlinear fashion, as a function of the depolarization in the post-synaptic cell. The proposal (Koch 1987) is therefore that the descending stimulation of the LGN causes long-lasting subthreshold depolarization. In this

proposal, the ascending stimulation involves receptors of the NMDA type. If they are activated while the units are still in a depolarized state, the response will be enhanced significantly. The long-lasting depolarization could be contributed by post-synaptic responses with slow time course, similar to the persistent Na^+ channel, or the I_T calcium channel (McCormick 1990). They can also be contributed by the activation of distal parts of the dendritic tree. Simulations of pyramidal cells (Stratford *et al.* 1989) have shown that such stimulation can have a significant temporal extent. It is also of interest to note in this regard that McGuire *et al.* (1984) have found that the connections from layer 6 to layer 4 in the cat (assumed in the model to be of the priming type) make connections to distal, slender branches of the dendritic tree.

The above two mechanisms could serve, individually or in combination, to implement connections of the priming type. Other mechanisms, not considered here, might participate as well. Although the details are not known, it appears that synaptic mechanisms for priming connections are physiologically plausible.

Modifiable connections

Several of the excitatory connections described above may be modifiable by experience. The learning and reinforcement of sequences discussed in Part I suggests, in particular, that the connections between successive patterns in a stream should be reinforced following the successful establishment of a complete sequence.

Two predictions follow. First, the synaptic connections between successive stages in each stream, such as AS $\rightarrow$ 4, or 6 $\rightarrow$ DS, should be modifiable. Second, the synaptic reinforcement should be particularly effective if the stimulation along the stream is preceded by a stimulation of the complementary stream. Long-term reinforcement effects have been reported in the visual cortex (Teyler *et al.* 1990), but they are sometimes difficult to elicit (Artola & Singer 1987). It would be of interest to test whether some of these effects can be facilitated by an appropriate stimulation of the complementary stream.

Finally, some of the reported LTP effects in the neocortex, including those that persist for only intermediate time intervals (minutes to a few hours) rather than very long ones, could be useful in implementing context

effects of the type discussed in Part I.

The feedback projections

One of the most characteristic properties of the visual (as well as non-visual, Felleman & Van Essen 1991) cortex is the reciprocity of the connectivity pattern. As a general rule, whenever area A sends a forward projection to area B, there is also a back-projection from B to A. In the sequence-seeking scheme, this rule follows from the counter-streams architecture (Fig. 3). The general rule in this scheme is that every forward connection A $\rightarrow$ B is accompanied by a reciprocal connection $\bar{B} \rightarrow \bar{A}$.

In the case of the projection from the LGN to the primary visual cortex, there is a possible reason why the back-projection may in fact be considerably more extensive than the forward projection. The reason is related to Marr's theory of the cortex (Marr 1970, 1971). In his model, Marr found that it is important to ensure that the input to the cortex will be sparse; that is, that the number of active fibers in the population encoding of a pattern will be kept relatively small. (A similar conclusion was reached by Willshaw *et al.*, 1969.) This consideration implies that at the input stage from the LGN to the cortex it is desirable to re-encode the input in a sparser manner (this may be one of the functions of the V1 granular layer). The result will be a mismatch between the LGN and V1: patterns in V1 will be represented by a larger number of units, and the number of descending connections will be correspondingly larger.

Effects of the feedback projection

According to the sequence-seeking scheme, the physiological effects of the descending projections are expected to be of two distinct types. The first effect is a priming, or modulatory effect. After a descending stream has been activated, it leaves the corresponding ascending stream in a primed state, making it more responsive to visual stimulation. The second effect is direct activation: a pattern on the descending path in a lower area can be directly activated by a higher-level pattern in the stream.

Both of these effects of back-projections have in fact been observed in physiological studies. Modulatory effects have been described by Nault *et al.* (1990, also Sandell & Schiller 1982) who studied the effects on area 17

of reversibly inactivating area 18 in the cat. The main effect found was a reduced responsiveness of area 17 units to visual stimulation when area 18 was inactivated. Similar modulatory effects have been found in the LGN, induced by cortical back-projections (e.g., Geisert *et al.* 1981, Marrocco *et al.* 1982). The reported effects were relatively subtle, and were both excitatory and inhibitory.

A prediction of the model regarding the modulatory effects is that facilitation by the back projections will not require strict temporal coincidence: a descending activation should be able to facilitate an ascending one even if it precedes it in time. Another prediction of the model is that similar modulatory effects are also likely to be induced by ascending signals on descending ones.

The priming effects are expected in the model to be subtle, and they could lead to both excitatory and inhibitory effects. An activation of a descending sequence will prime the complementary patterns. At the same time, due to competition between patterns and inhibition among them as discussed above, the effect on other units could be inhibitory. At the level of the LGN, the effects are expected to be mainly, (perhaps exclusively), of the modulatory type.

A direct excitatory effect of the descending projections was found by Miguard and Malpeli (1991). They have blocked the direct ascending input from the LGN to V1 in the cat, and found that layers in the upper layers of V1 could still be strongly activated via the back-projections from V2. They conclude that their findings do not support the modulatory view of the feedback circuits, since the back-projections can in fact drive V1 units vigorously. From the point of view of the sequence-seeking model, both of these interactions are in fact expected. Furthermore, it is expected that they will be segregated into two distinct sub-populations, one on the ascending, the other on the descending stream. Referring back to Fig. 2, pattern $\bar{B}$ on the descending stream can be directly driven from higher patterns along the descending stream. In contrast, pattern B on the ascending stream is expected to show priming, modulatory effects of the descending stream.

Due to the direct connections on the descending stream, under the appropriate conditions one may expect to see complex patterns of activity induced by the descending stream even in lower-level visual areas. These

patterns will change with time, as sequences that have been explored decay and are replaced by new ones. Such patterns may be difficult to observe with single-cell recordings, but they may be observable with suitable optical recording techniques with voltage-sensitive dyes.

Conclusion

In the sequence-seeking model, the basic operation of the neocortex is to seek sequences of transformations, or mappings, linking source and target representations. This operation is implemented in a counter-streams structure, in which multiple 'bottom up' and 'top down' sequences are explored in parallel along separate ascending and descending pathways. The individual mappings and transformations will differ in different cortical areas, but the basic flow of the process, the parallel bi-directional exploration of a large number of sequences, is assumed to be common to different parts of the cortex.

Many elaborations and local variations may exist, but the basic scheme is uniform and in principle straightforward. Key properties that are expected to contribute to the scheme's power include: the simultaneous exploration of a large number of alternative sequences, the relative simple, uniform, and extensible structure, the flexible use of 'bottom up' and 'top down' sequences that can meet at any level, the broad and indirect context effects, and the learning of complete sequences by a simple local reinforcement rule.

Acknowledgement: I would like to thank C. Gilbert, E. Grimson, C. Koch, J. Maunsell, D. Mumford, T. Poggio, and W. Richards for their helpful comments.

APPENDIX

Multiple Co-Existing Patterns

In the sequence-seeking scheme it is desirable to explore in parallel a large number of different sequences. In an actual implementation, there will be constraints on the maximal possible number of co-active patterns. In particular, if patterns are represented by the activity of different sub-

populations of neurons, the co-activation of too many patterns will lead to interference between patterns and possible confusion.

In this appendix, an approximate bound on the number of co-activated patterns is derived in the context of a simple associative memory model, of the type described by Willshaw, Baneman, and Longuet-Higgins (1969).

In this model, the memory consists of n input and n output neurons, with complete interconnections. Pairs of patterns $(a_i - b_i)$ can be written into this memory. This means that if any of the patterns a_i is presented as an input to the system, it will reproduce the corresponding b_i (even if a_i is partially corrupted). Willshaw *et al.* have shown that under certain assumptions, their model can store about $ln2 \times n$ different pattern pairs.

For the co-activation of patterns, the relevant question is what happens if a number of different patterns $(a_1, a_2, \ldots, a_k)$ are presented to the system simultaneously. We would like to get at the system's output all of the patterns $(b_1, b_2, \ldots, b_k)$ but no other patterns. The question is, essentially, how large can k be without leading to possible confusion?

The patterns in this model are composed of 1's ("active" units) and 0's ("inactive" ones). Assume that k patterns, each composed of m 1's (and $n - m$ 0's) are presented simultaneously. The total number of 1's in the input can be estimated as follows. The probability of a given input line not being activated by a given pattern is $1 - \frac{m}{n}$. The probability of not being activated by k independent inputs is (for $m << n$):

$$(1 - \frac{m}{n})^k \approx e^{-\frac{mk}{n}} \qquad (A1)$$

and the number M of active input units is therefore

$$M = n(1 - e^{-\frac{mk}{n}}). \qquad (A2)$$

How large can M be without causing possible problems? We would like the probability of an output pattern b being activated by mistake by the input to be small. A pattern b will be activated by chance if all (or almost all) of its input will happen to be activated. The probability that this will happen for a given pattern b is $(M/n)^m$. To keep the average number of such errors small, we will require: $n \times (M/n)^m < 1$ (since there are close to n patterns stored). For optimal use of the associative memory, Willshaw

et al. have determined that m should be about $log_2 n$. This leads to the requirement:

$$M < 0.5n \qquad (A3).$$

Combining this with (A2):

$$M = n(1 - e^{-\frac{m}{n}k}) = 0.5n \qquad (A4)$$

therefore:

$$k = \frac{ln2 \times n}{m}. \qquad (A5)$$

Or, assuming that $m = log_2 n$:

$$k = \frac{ln2 \times n}{log_2 n}. \qquad (A6)$$

In this model, then, the maximal number of patterns that can be stored is $P = ln2 \times n$, and the maximal number of co-active patterns should not exceed $\frac{P}{log_2 n}$.

As expected, the fact that the patterns are not represented independently places constraints on the number of patterns that can be activated simultaneously. At the same time (at least within this associative memory model), the number of co-active patterns can still be large.

An additional constraint worth noting is where the mapping from one set of patterns to the next may not be invertible. For example, suppose that patterns are being classified into "symmetric" vs. "non-symmetric." Given a pattern, the "forward" mapping (i.e., classifying it as symmetric / non-symmetric) can be computed. In the opposite direction, however, knowing that a pattern is symmetric is insufficient for reconstructing the pattern. In the sequence-seeking scheme, it was assumed that if a transformation $A \rightarrow B$ can take place in the ascending stream, then the transformation $\bar{B} \rightarrow \bar{A}$ can take place in the descending stream. In the example above, the situation is asymmetric: $A \rightarrow B$ is well-defined, while the inverse, $B \rightarrow A$, is not. The main implication of this constraint is to require certain temporal ordering on the expansion of sequences, e.g., that $A \rightarrow B$ on the ascending stream will have to occur before $\bar{B} \rightarrow \bar{A}$ on the descending stream.

REFERENCES

Albowitz, B & Kuhnt, U. 1991. Optical recordings of evoked activity in guinea pig neocortical slices. *Soc. Neurosci. Abstracts*, **17**, 114.

Albus, J.S. 1971. A theory of cerebellar functions. *Math.Bio.*, **10**, 25-61.

Artola, A. and Singer, W. Long-term potentiation and NMDA receptors in visual cortex. 1987. *Nature,* **330**, 649-652.

Amit, D.J. 1989. **Modeling Brain Function**. Cambridge, UK: Cambridge University Press.

Andersen, R.A., Asanuma, C., Essick, G., and Siegel, R.M. 1990. Cortico-cortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J. Comp. Neurol.*, **296**, 65-113.

Barlow, H.B. 1985. Cerebral cortex as model builder. In: **Models of the Visual Cortex**, D. Rose and V.G. Dobson (eds). New York: John Wiley and Sons, Ltd.

Blasdel, G.G. and Lund, J.S. 1983. Termination of afferent axons in macaque striate cortex. *J. Neurosci,* **3**, 1389-1413.

Boussaud, D., Ungerleider, L.C., and Desimone, R. 1990. Pathways for motion analysis: Cortical connections of the medial superior temporal and fundus of the superior temporal visual areas in the macaque. *J. Comp. Neurol.*, **296**, 462-495.

Brady, R.M. 1985. Optimization processes gleaned from biological evolution. *Nature,* **317**, 804-806.

Carpenter, G., and Grossberg, S. 1987. A massively parallel architecture for a self-organizing neural pattern recognition machine. *Comp. Vis., Graphics and Image Proc.*, **37**, 54-115.

Creutzfeldt, O.D. 1978. The neocortical link: Thoughts on the generality of structure and function of the neocortex. In: **Architectonics of the Cerebral Cortex**, M.A.B. Brazier and H. Petsche (eds). New York: Raven Press.

Crick, F. and Asanuma, C. 1986. Certain aspect of the anatomy and physiology of the cerebral cortex. In: J.L. McClelland and D.E. Rumelhart (eds.), **Parallel Distributed Processing** (Vol. 2), Cambridge: MIT Press.

Desimone, R., Li, L., Lehky, S. Ungerleider, L.G., and Mishkin, M. 1990. Effects of V4 lesions on visual discrimination performance and on responses of neurons in inferior temporal cortex. *Soc. Neurosci. Abstracts*, **16**, 621.

Edelman, G.M. 1978. Group selection and phasic re-entrant signalling: A theory of higher brain functions. In: G.M. Edelman and V.B. Mountcastle, **The Mindful Brain**, Cambridge, MA: MIT Press.

Esguerra, M., Kwon, Y.H., and Sur, M. 1989. NMDA and non-NMDA receptors mediate retinogeniculate transmission in cat and ferret LGN *in vitro*. *Soc. Neurosci. Abstracts*, **15**, 175.

Esguerra, M., and Sur, M. 1990. Corticogeniculate feedback gates retinogeniculate transmission by activating NMDA receptors. *Soc. Neurosci. Abstracts*, **16**, 159.

Felleman, D.J., and Van Essen, D.C. 1991. Distributed hierarchical processing in primate visual cortex. *Cerebral Cortex*, in press.

Ferster, D., and Lindström, S. 1985. Augmenting responses evoked in area 17 of the cat by intracortical axon collaterals of cortico-geniculate cells. *J.Physiol.*, **367**, 217-232.

Ferster, D., and Lindström, S. 1985. Synaptic excitation of neurones in area 17 of the cat by intracortical axon collaterals of cortico-geniculate cells. *J.Physiol.*, **367**, 233-252.

Fitzpatrick, D. Itoh, K and Diamond, I.T. 1983. The laminar organization of the lateral geniculate body and the striate cortex in the squirrel monkey (Samiri Sciureus). *J. Neurosci.* **3(4)**, 673-702.

Fitzpatrick, D. and Einstein, G. 1989. Laminar distribution and morphology of area 17 neurons projecting to the lateral geniculate nucleus in the macaque. *Soc. Neurosci. Abstracts*, **15**, 1398.

Friedman, D.P. 1983. Laminar patterns of termination of corticocortical afferents in the somatosensory system. *Brain Res.*, **273**, 147-151.

Geisert, E.E., Langsetmo, A., and Spear, P.D. 1981. Influence of the corticogeniculate pathway on response properties of cat lateral geniculate neurons. *Brain Res.*, **208**, 409-415.

Grossberg, S. (Ed.) 1988. **Neural Networkd and Natural Intelligence.** Cambridge, MA: M.I.T. Press.

Grieve, K.L., Murphy, P.C., and Sillito, A.M. 1991. Inhibitory and excitatory components to the subcortical and cortical influence of layer VI cells in the cat visual cortex. *Soc. Neurosci. Abstracts*, **17**, 629.

Holland, J.H. 1975. **Adaptation in Natural and Artificial Systems.** Ann Arbor: Univ. of Michigan Press.

Hendrickson, A.E., Wilson, J.R. and Ogren, M.P. 1978. The neuroanatomical organization of pathways between the dorsal LGN and visual cortex in old world and new world primates. *J. Comp. Neurol.*, **182**, 123-136.

Hopfield, J.J. 1982. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA*, **79**, 2554-2558.

Karni, A, and Sagi, D. 1991. Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proc. Nat. Acad. Sci. USA,* **88**, 4966-4970.

Kasper, E., Larkman, A.U., Blakemore, C. and Judge, S. 1991. Physiology and morphology of identified projection neurons in rat visual cortex studied *in vitro*. *Soc. Neurosc. Abstracts,* **17**, 114.

Katz, L.C., Burkhalter, A. and Dreyer, W.J. 1984. Fluorescent latex microspheres as a retrograde neuronal marker for *in vivo* and *in vitro* studies of visual cortex. *Nature*, **310**, 498-500.

Kennedy, H. and Bullier, J. 1985. A double labeling investigation of the afferent connectivity to cortical areas V1 and V2 of the macaque monkey. *J. Neurosci.* **5**, 2815-2830.

Koch, C. 1987. The action of the corticofugal pathway on sensory thalamic nuclei: A hypothesis. *Neuroscience*, **23**, 399-406.

Livingston, M.S. and Hubel, D.H. 1982. Thalamic inputs to cytochrome oxidase-rich regions in monkey visual cortex. *Proc. Nat. Acad. Sci.*, **79**, 6098-6101.

Lund, J.S., Fitzpatrick, D. and Humphrey, A.L. 1985. The striate visual cortex of the tree shrew. In: Peters and E.J. Jones (eds), **Cerebral Cortex**, Vol. 3, A. New York: Plenum Press.

Lund, J.S. 1987. Local circuit neurons of macaque monkey striate cortex: I. Neurons of laminae 4C and 5A. *J. Comp. Neurol.*, **257**, 60-92.

Lund, J.S. 1988a. Excitatory and inhibitory circuitry and laminar mapping strategies in primary visual cortex of the monkey. In: G.M. Edelman, W.E. Gall, and W.M. Cowan (eds), **Signal and Sense: Local and Global Order in Perceptual Maps**, New York: John Wiley and Sons.

Lund, J.S. 1988b. Anatomical organization of macaque monkey striate visual cortex. *Ann. Rev. Neurosci.*, **11**, 253-288.

Marr, D. 1970. A theory for cerebral neocortex. *Proc. Roy. Soc. Lond. B.*, **176**, 161-234.

Marr, D. 1971. Simple memory: A theory for archicortex. *Phil. Trans. Roy. Soc. Lond. B.*, **262**, 23-81.

Marrocco, R.T., McClurkin, J.W., and Young, R.A. 1982. Modulations of lateral geniculate nucleus cell responsiveness by visual activation of the corticogeniculate pathway. *J. Neurosci.*, **2**, 256-263.

Martin, K.A.C. 1988. From singe cells to simple circuits in the cerebral cortex. *Quart. J. Exp. Physiol.*, **73**, 637-702.

Maunsell, J.H.R., and Van Essen, D.C. 1983. The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *J. Neurosci.*, **3**, 2563-2586.

McCormick, D.A. 1990. Membrane properties and neurotransmitter actions. In:f G.M. Shepard (ed.), **The Synaptic Organization of the Brain** , 3rd edition, New York: Oxford Univ. Press.

McGuire, B.A., Hornung, J.-P., and Gilbert, C.D. 1984. Patterns of synaptic input to layer 4 of cat striate cortex. *J. Neurosci.*, **4**, 3021-3033.

Mignard, M., and Malpeli, J.G. 1991. Paths of information flow through visual cortex. *Science*, **251**, 1249-1251.

Miller, K.D., Chapman, B., and Stryker, M.P. 1989. Visual responses in adult cat visual cortex depend on *N*-methyl-D-aspartate receptors. *Proc. Natl. Acad. Sci. USA*, **86**, 5183-5187.

Mitzdorf, U. and Singer, W. 1978. Prominent excitatory pathways in the cat visual cortex (A17 and A18): A current source density analysis of electrically evoked potentials. *Exp. Brain Res.,* **33,** 371-393.

Mitzdorf, U. and Singer, W. 1979. Excitatory synaptic ensemble properties in the visual cortex of the macaque monkey: A current source density analysis of electrically evoked potentials. *J. Comp. Neur,* **187**, 71-84.

Mumford, D. 1990. On the computational architecture of the neocortex. I. The role of the thalamo-cortical loop. *Biol. Cyber.* **65**(2), 135-145.

Nault, B., Michaud, Y., Morin, C., Casanova, C., and Molotchnikoff, S. 1990. Responsiveness of cells in area 17 after local interception of the descending path from area 18. *Soc. Neurosci. Abstracts*, **16**, 1219.

Nauta, W.J.H. and Feirtag, M. 1986. **Fundamental Neuroanatomy** New York: W.H. Freeman.

Pallas, S.L. Roe, A.W. and Sur, M. 1990. Visual projections induced into the auditory pathway of ferrets. I. Novel inputs to primary auditory cortex (AI) from the LP/Pulvinar complex and the topography of the MGN-AI projection. *J. Comp. Neurol.,* **298** 50-68.

Peterson, C. 1990. Parallel distributed approaches to combinatorial optimization: Benchmark studies on traveling salesman problem. *Neural Computation,* **2**, 261-269.

Poggio, T. 1990. A Theory of How the Brain Might Work. *Artificial Intelligence Laboratory Memo #1253.* Cambridge, MA: Massachusetts Institute of Technology.

Posner, M.I., Snyder, C.R.R., and Davidson, B.J. 1980. Attention and the detection of signals. *J. Exp. Psychol., General,* **109**, 160-174.

Reeke, G.N., Finkel, L.H., Sporns, O. and Edelman, G.M. 1990. Synthetic neural modeling: A multilevel approach to the analysis of brain complexity. In: G.M. Edelman, W.E. Gall, and W.M. Cowan (eds), **Signal and Sense: Local and Global Order in Perceptual Maps**, New York: John Wiley and Sons.

Rockel, A.J., Hiorns, R.W. and Powell, T.P.S. 1980. The basic uniformity in structure of the neocortex. *Brain,* **103**, 221-244.

Rockland, K.S., and Lund, J.S., 1983. Instrinsic laminar lattice connections in primate visual cortex. *J. Comp. Neurol.*, **216**, 303-318.

Rockland, K.S., and Pandya, D.N. 1979. Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Res.*, **179**, 3-20.

Rockland, K.S., and Virga, A. 1989. Terminal arbors of indivual "feedback" axons projecting from area V2 to V1 in the macaque monkey: A study using immunohistochemistry of anterogradely transported *Phaseolus vulgaris*-leucoagglutinin. *J .Comp. Neurol.*, **285**, 54-72.

Roe, A.W. 1991. Functional transformations of visual input by auditory thalamus and cortex: An experimentally induced visual pathway in ferrets. Ph.D. Thesis, M.I.T., Department of Brain and Cognitive Science.

Roe, A.W. Hahm, J. and Sur, M. 1991. Experimentally induced establishment of visual topography in auditory thalamus. *Soc. Neurosc. Abstracts,* **17**, 898.

Sandell, J.H., and Schiller, P.H. 1982. Effect of cooling area 18 on striate cortex cells in the squirrel monkey. *J. Neurophysiol.*, **48**, 38-48.

Schiller, P.H. and Lee, K. 1991. The role of the primate extrastriate area V4 in vision. *Science,* **251**, 1251-1253.

Sejnowski, T.J. 1986. Open questions about computation in cerebral cortex. In: J.L. McClelland and D.E. Rumelhart (eds.), **Parallel Distributed Processing** (Vol. 2), Cambridge: MIT Press.

Sherman, S.M., Scharfman, H.E., Lu, S.M., Guide, W., and Adams, P.R. 1990. *N*-methyl-D-aspartate (NMDA) and non-NMDA receptors participate in EPSPs of cat lateral geniculate neurons recorded in thalamic slices. *Soc. Neurosci. Abstracts*, **16**, 159.

Silva, L.R., Amitai, Y. and Connors, B.W. 1991. Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. *Science,* **251**, 432-435

Stratford, K., Mason, A., Larkman, A., Major, G., and Jack, J. 1989. The modeling of pyramidal neurones in the visual cortex. In: **The Computing Neuron**, R. Durbin, C. Miall, and G. Mitchison (eds). Reading, MA: Addison-Wesley.

Sur, M. Garraghty P.E. and Roe, A.W. 1988. Experimentally induced visual projections into auditory thalamus and cortex. *Science,* **242** 1437-1441.

Teyler, T., Aroniadou, V., Berry, R.L., Barroni, A. DiScenna, P, Grover, L. and Lambert, N. 1990. LTP in neocortex. *Seminars in the Neuroscience,* **2**, 365-379.

Ullman, S. 1984. Visual routines. *Cognition,* **18**, 97-159.

Van Essen, D.C. 1985. Functional organization of primate visual cortex. In: **Cerebral Cortex**, Vol. 3, A. Peters and E.J. Jones (eds). New York: Plenum Press.

White, E.L. 1989. **Cortical Circuits**. Boston: Birkhäuser.

Willshaw, D.J., Buneman, O.P., and Longuet-Higgins, H.C. 1969. Non-holographic associative memory. *Nature*, **222**, 960-962.

Zeki, S. and Shipp, S. 1988. The functional logic of cortical connections. *Nature,* **335**, 311-317.