Feature Article

Sequence Seeking and Counter Streams: A Computational Model for Bidirectional Information Flow in the Visual Cortex

Shimon Ullman

Department of Brain and Cognitive Science, MIT, Cambridge, Massachusetts and Department of Applied Mathematics, The Weizmann Institute, Rehovot, Israel 76100

A computational model is proposed for some general aspects of information flow in the visual cortex. The basic process, called "sequence seeking," is a search for a sequence of mappings, or transformations, linking source and target patterns. The process has two main characteristics: it is bidirectional, bottom-up as well as top-down, and it explores in parallel a large number of alternative sequences. This operation is performed in a "counter streams" structure, in which multiple sequences are explored along two complementary pathways, an ascending and a descending one, seeking to meet. A biological embodiment of this model in cortical circuitry is proposed. The model serves to account for known aspects of cortical interconnections and to derive new predictions.

This article describes a computational model for some general aspects of information flow in the visual cortex. The model combines computational and psychophysical considerations with data concerning connectivity patterns in the visual cortex, primarily interlaminar and connections between different cortical areas.

The focus of the model is on vision (primarily the task of visual recognition) and the visual cortex. The proposed computation has, however, some useful generic aspects, and the possible applicability of the scheme to other domains is briefly considered. The first part of this report outlines the proposed computation, termed "sequence seeking," and the second part outlines its biological embodiment in a "counter streams" structure. The model is used to account for known features of cortical circuitry, and to derive a number of new predictions.

Sequence Seeking and Counter Streams

A general task frequently faced by the brain is one of establishing a link between two different representations. For example, in visual recognition, the task involves establishing a connection between an incoming pattern and stored object representations in visual memory. The two will often fail to match exactly, due to changes in size, position, viewing direction, and so on. The problem is therefore not merely one of direct pattern matching; considerable processing needs to be performed to overcome the possible differences between a given image and previously stored patterns. A common view is, therefore, that prior to the matching the input is processed through a sequence of stages that include, for example, edge detection, extracting features of varying complexity, normalization for size, position, orientation, and so on. The model below modifies and extends this view based on the use of two general strategies that are supported by computational and psychological considerations. The first is to employ a bidirectional search, where the matching of patterns can occur at intermediate levels rather than some "topmost" level. Second, rather than following a single path, multiple processing alternatives are explored in parallel. The next two sections describe these general strategies, followed by a proposed model that combines bidirectional processing with the parallel exploration of multiple alternatives.

Bidirectional Processing

In performing visual recognition, an input pattern P is matched with an appropriate stored representation, from a set of stored models M_i. The processing stages required to obtain the match of P with the appropriate stored model can be described in terms of applying a sequence of processes, or transformations, that compensate for possible differences between the two, for example, in scale, position, viewing direction, illuminations conditions, and so on (Ullman, 1993). These processing stages could be applied to the input P, the models M_i, or both. Some of the processes are more naturally applied to the incoming pattern. For example, a simple transformation, such as overall shift or scaling, is best applied to the input pattern, because then it will be applied to a single pattern, rather than to multiple stored models. However, other transformations are specific to a stored model, for example, how a given 3D object appears from different viewing directions, or how a face may transform by facial expressions. Such object-specific transformations are more difficult to compensate for by processing the input pattern alone, but they can be handled effectively by processes that utilize model-specific information, accumulated through past experience, for example, about the object's 3D shape or the distortions it may undergo (Lowe, 1985; Ullman, 1989; Grimson, 1990; Yuille and Hallinan, 1992). An attractive overall strategy is therefore to apply a bidirectional process, where processing can be applied simultaneously to both the input and the stored patterns. An additional advantage is that a bidirectional computation can also be considerably more economical than unidirectional processing in terms of the number of patterns explored during the computation, as we shall see further below in discussing simulation results.

There is considerable computational and psychophysical evidence in support of the view that visual recognition is indeed not a unidirectional process, but requires a cooperation of processes applied to both the input image and stored object models. The evidence points to a combination of bottom-up and top-down processing, where "bottom-up" processes are involved with the analysis of the incoming image, and "top-down" processes originate with stored models and information associated with them. In computer vision, the integration of bottom-up with top-down processing has been a major concern. During the 1970s, the emphasis was placed heavily on top-down processes. Fundamental difficulties with building computer vision systems led to the view that the processing must be guided primarily by knowledge associated with stored models of objects and scenes, and systems in the 1970s were constructed using this approach (Freuder, 1974; Tenenbaum and Barrow, 1976). Following in part the work of Marr (1982), the emphasis shifted toward bottom-up processing, but it was also evident that a key issue is the integration of bottom-up with top-down processing (Marr, 1982; Grimson, 1990), and both processes are used in current recognition systems (e.g., Lowe, 1985; Ullman, 1989; Grimson, 1990; Yuille and Hallinan, 1992). Psychological studies have supported this view, and have shown the importance of both types of processes in object recognition, segmentation, and scene perception (Palmer, 1975; Potter, 1975; Biederman et al., 1982), as will be described in more detail in subsequent sections.

In the sequence-seeking scheme presented below bottom-up and top-down processes are equally important. The scheme uses two streams of processing, an ascending one starting at the input, and a descending one starting at the stored models. Their integration is achieved by the interactions between two complementary processing streams. From a biological standpoint, these will correspond to the "forward" and "backward" connections between cortical areas. A key ingredient of the counter-streams scheme is the suggestion that, roughly speaking, the ascending pathways subserve mainly bottom-up processing, and the descending pathways mainly top-down processing, and certain interactions between the streams provide a mechanism for integrating together the two types of processing.

Exploring Multiple Alternatives

A second general strategy is designed to face the possibility that 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 efficient computation, it will be necessary to explore simultaneously multiple alternatives. For example, to deal with size variations, the input will be processed at multiple scales in parallel. Similarly, in the case of recognition, the match of an input image and stored model will be attempted by the system at different orientations, positions, and 3D viewing directions. In many models of visual processing, the input pattern undergoes a single sequence of processing stages. In contrast, in the sequence-seeking scheme an input pattern gives rise to multiple sequences of transformations and mappings that are explored in parallel. The different transformations and mappings explored by the system should be taken here in a broad sense: in addition to dealing with geometric transformations such as changes in size, position, and orientation, the processing may involve the recovery of different properties such as color, motion, texture, and 3D shape, as well as exploring alternative ways of representing the pattern, for example, in terms of its parts and its abstract shape properties (Ullman, 1989).

A simplified example can help to illustrate the processing. Suppose that we attempt to recognize a face we have seen before, but under novel viewing conditions, in terms of size and 3D viewing direction. Such variations in viewing conditions can be handled by combining bidirectional search with the exploration of multiple alternatives. The top-down processes will generate internally several representations, corresponding to different 3D viewing directions, and the bottom-up processes will analyze the input at multiple scales in parallel. The system will explore the different alternatives for a possible match between the transformed versions of the input and stored model, and will be able to obtain in this manner the appropriate match, despite the initial discrepancy between them.

The need to evaluate multiple competing alternatives arises frequently in computer vision systems performing tasks such as segmentation and recognition. For example, some of the best-performing visual recognition systems (Lowe, 1985; Grimson, 1990) search for the best match between an input shape and a candidate internal model by exploring and comparing multiple (e.g., in the hundreds) possible 3D poses of the internal model. In current computer systems the multiple comparisons are performed sequentially. In a biological system, which is slower but inherently parallel, a better approach is to explore simultaneously as many alternatives as possible.

It is interesting to note that in several recent neural network models in areas such as arm control (Jordan and Jacobs, 1993) or handwriting recognition (Matan et al., 1992), the explicit exploration of multiple alternatives (e.g., by separate "expert subnetworks") proved useful in dealing with complex problems that were not handled effectively by more standard models. For example, in a network developed at AT&T for reading handwritten numerals (Matan et al., 1992), the letter identification stage is preceded by a process that segments the string into individual characters. This turned out to be a difficult task, and the segmentation was often unreliable. The approach adopted was to perform explicitly a number of different candidate segmentations, and later select the best solution among the different alternatives. The network also uses multiple-scale analysis: it processes the same input through separate mechanisms tuned to different scales of the letters, and again selects at the end the best alternative. This design proved to perform better than a single network trained to become scale invariant.

Another example is a network developed recently for controlling the 3D movements of a simplified arm (Jordan and Jacobs, 1993). In this work, the training of standard neural networks failed to converge to a satisfactory solution to the control problem. An approach that proved successful was to train a number of different subnetworks, each one capable of providing a good solution under restricted conditions. The combined problem is then treated by letting the different networks work on the problem individually, followed by a gating and selection stage.

Linking the Ascending and Descending Streams, the Counter-Streams Structure

The scheme proposed for combining bidirectional and multipath processing is diagrammed schematically in Figure 1*A*. The basic operation in this scheme is to seek a sequence of processing steps linking a pattern of activity (S in Fig. 1*A*) in one cortical area with stored representations (such as M_1, M_2) in another. The pattern S may arise from the image of an object, for example, a familiar face, and the patterns M_i represent stored object models, perhaps in visual area IT. The recognition of the viewed object involves multiple processing stages, applied in part to the incoming pattern and in part to the stored models, in an attempt to establish a match between the incoming pattern and a stored model. For example, the bottom-up stream may include processes that compensate for size and position variations, and the top-down stream processes compensating for illumination conditions and viewing direction. Intermediate patterns in the net will correspond to different representations of the object, for example, at different 3D orientations and scales. Biologically, the nodes in this schematic figure represent patterns of activity, for example, subpopulations of neurons acting together, possibly with some degree of synchrony (Abeles, 1991; Engel et al., 1992), and the arrows indicate how patterns activate subsequent patterns, for example, S can activate A_2, A_3, and A_4. Since different patterns may share neurons, implementation constraints will place some limitations on the coactivation of patterns; for example, patterns (B_2, B_3, B_4) may be prohibited from being coactive. In expanding the sequences down from M_1, only a subset of these patterns will be activated initially, and will later decay and be replaced by others.

The search is bidirectional, and a linking sequence is successfully established when the two streams of activation meet somewhere in this large network of interconnected patterns. How can a successful link of patterns be found in the system? The proposed scheme (Fig. 1*B*) 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 is assumed to leave behind

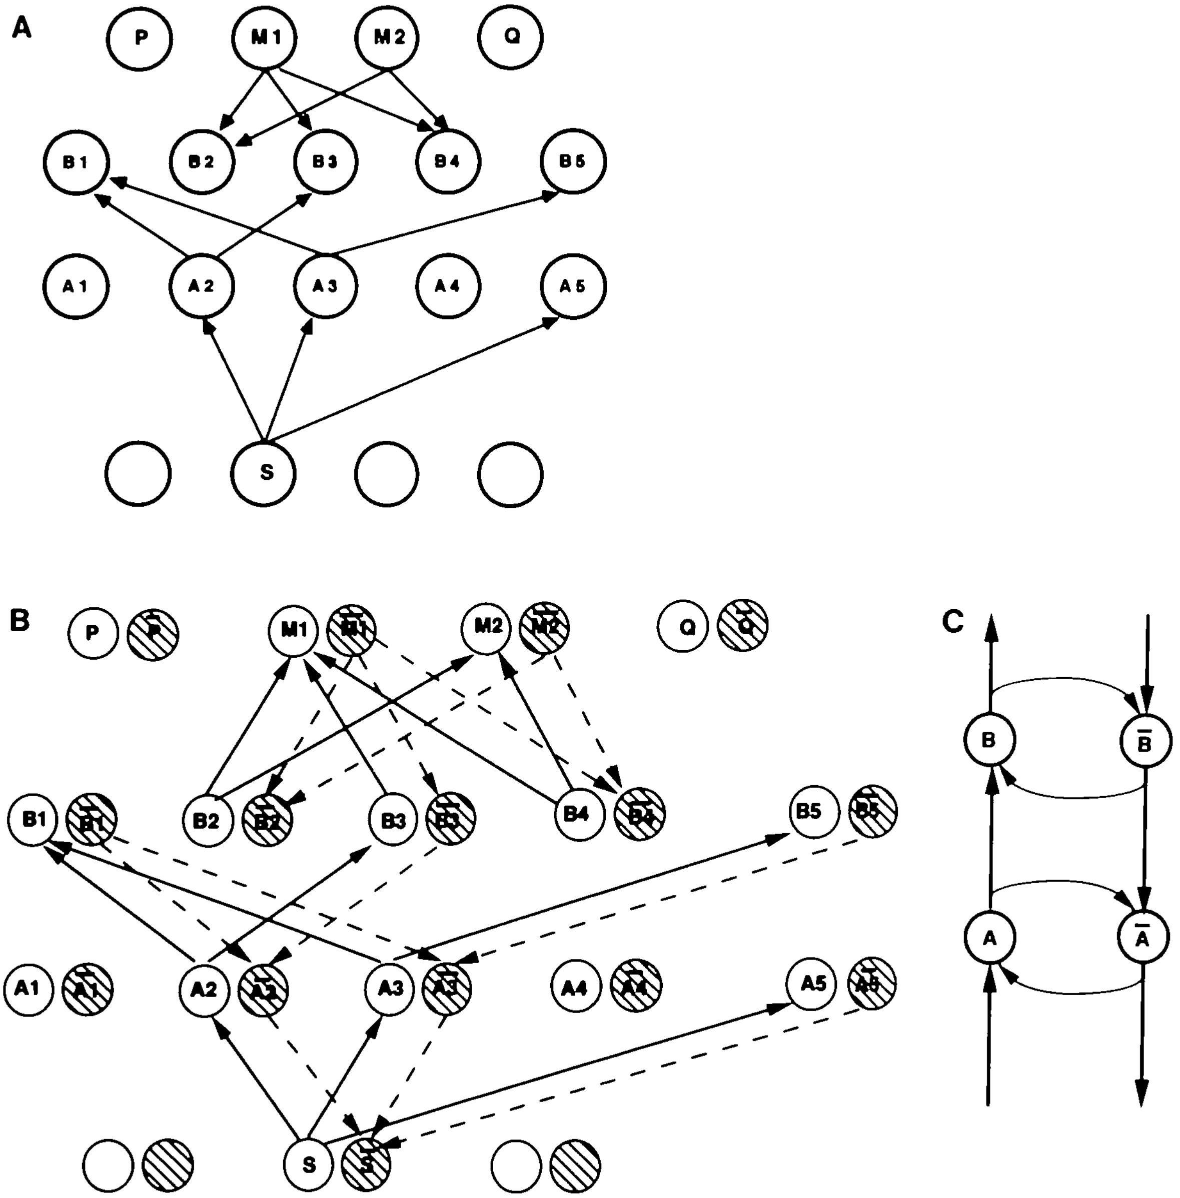

Figure 1. A, The sequence-seeking computation seeks a sequence of mappings linking a source pattern (*S*) in one area with stored representations (*M1, M2*) in another. *Nodes* represent patterns of activity (coactive populations of neurons); *arrows* indicate how patterns activate subsequent patterns. In expanding sequences only a subset of patterns will be activated initially, and will later decay and be replaced by others. The processing is bidirectional, and a linking sequence is successfully established when the two searches meet somewhere in a large network of interconnected patterns. B, Similar to *A*, except that each node is split into two complementary ones. The ascending and descending streams proceed along complementary pathways. When a track is being traversed in one stream, it leaves behind a primed trace in the complementary stream. C, The basic unit of the counter-streams structure: Patterns *A, B* on the ascending, $\bar{A}$, $\bar{B}$ on the descending path. *Thin arrows* denote connections of the priming type. This repeating unit is embedded in a network of richly interconnected patterns.

a primed trace in the complementary stream, making it more readily excitable, as explained further below. The scheme shown schematically in Figure 1*B* is similar to Figure 1*A*, except that each node is now split into two complementary nodes (populations of neurons), for example, B_2 in Figure 1*A* is now split into B_2 on the ascending pathway and its complementary pattern $\bar{B}_2$ on the descending one.

The full bidirectional 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 downward along the descending pathway. (We will see below how some models, such as M_1, M_2, can be selected from a larger population of stored patterns.) Whenever a track (subsequence) is being traversed on either stream, the complementary track remains in a primed state, ready to be activated. Not all of the possible sequences are expanded simultaneously, and already-primed patterns are activated with priority. The result will be a mechanism that searches for linking sequences in the network. 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. This is an example of a linking taking place between the two streams; that is, a node, or group of neurons, on the ascending streams (A_2 in the example) and its counterpart ($\bar{A}_2$) on the descending stream, have both been activated within a limited time interval (up to a few hundred milliseconds in the case of typical recognition).

The activity will then proceed along the primed traces, since, for example, A_2 will activate next B_3 (which is primed) rather than alternative, nonprimed nodes. This will result, therefore, in the immediate activation of the complete sequences $S \to M_1$ and $\bar{M}_1 \to \bar{S}$, establishing a complete link between the source and target patterns. This will also select M_1 as the stored pattern corresponding to the input image S, thereby serving to recognize S as an instance of M_1. In this manner, as a result of the priming, the top-down processing guides and "paves the way" for the bottom-up processing. The task of relating the sensory input to the appropriate stored representations is achieved by the cooperation of bottom-up and top-down processing, and by exploring, in each direction, multiple alternatives.

The linking process described above achieves a number of goals. First, a link between the ascending and descending streams can take place at any intermediate level; this has the advantage that the overall task can be split in a flexible manner between bottom-up and top-down processing. Second, to establish a link, the ascending and descending patterns need not arrive at a given 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. Third, the proposed separation between the streams avoids possible intermixing between data supported by the input and states explored internally by the system. This distinction is crucial in a system that uses both top-down and bottom-up processes. For example, top-down processing can initiate the activation of an internal model (such as $\bar{M}_1$ in Fig. 1*B*), but this event will be distinguished in the system from the activation of an internal model (M_1 in Fig. 1*B*) on the ascending stream by a sequence originating at the sensory input.

In terms of connectivity, the excitatory connections between patterns are predominantly reciprocal, obeying the following general rule (Fig. 1*C*): whenever A is connected to B, there is a back-connection from $\bar{B}$ to $\bar{A}$, with cross-connections between A and $\bar{A}$ and B and $\bar{B}$. (Inhibitory connections also play a role, but will not be discussed.) The cross-connections are assumed to have a priming effect; when B, for instance, is activated, it also provides input to $\bar{B}$, making it more readily excitable by a subsequent input along the descending stream. The reciprocity of the connections is an inherent aspect of the model, and it is also a distinguishing feature of cortical connectivity (although some exceptions have been noted; see Distler et al., 1991; Rockland et al., 1992). It should also be noted that although the counter-streams structure uses "forward" and "backward" connections, it does not necessarily imply a hierarchical structure; it can incorporate a more general structure as long as the above connectivity rule is obeyed.

In summary, the sequence-seeking process above has two main characteristics: it is bidirectional, and it explores multiple alternatives simultaneously. The basic structure of the model is relatively straightforward, comprising two 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 descending pathways are used in the model as the anatomical substrate for top-down processing. This role can be contrasted with other models, where the descending projections are used for different purposes: controlling selective attention (Fukushima, 1986; Koch, 1987), grouping and figure-ground segregation (Okajima, 1991; Sporns et al., 1991), learning processes (Zipser and Rumelhart, 1990), modulating cortical output to other visual centers (Sandell and Schiller, 1982) or to correlate and synchronize the activity of interrelated neuronal groups (Edelman, 1978; Tononi et al., 1992).

Compared with other models (e.g., Geman and Geman, 1984; Poggio et al., 1985; Mumford, 1992), the scheme places emphasis on the parallel exploration and selection of multiple alternatives, rather than relaxation and iterative computations. Timing considerations (Maunsell and Gibson, 1992; Rolls et al., 1991; Thorpe et al., 1991) appear to place rather stringent restrictions on the use of multi-iteration relaxation processes in tasks such as visual recognition. A visual cortical area may introduce an average delay of about 10-15 msec, and there are several (about six) stations spanning the hierarchy from V1 to anterior IT. This suggests that visual processing should usually require a limited number of sweeps through the system. It is desirable, therefore, especially for an inherently parallel system, to explore multiple alternatives simultaneously, rather than explore and refine them in sequence.

The bidirectional processing above provides the skeleton of the computation, a number of elaborations and properties of the basic process are discussed below.

Express Lines

The bidirectional process raises an important question regarding the activation of stored models for top-down processing. To initiate appropriate top-down processing, some initial selection, and subsequent refinement, of a relevant subset of stored models is required. The next two sections propose two mechanisms for this task, both supported by evidence regarding human perception.

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 yet be possible to identify the individual face. It would be advantageous under these circumstances to expand preferentially face-related sequences, possibly at the expense of others. Such an effect can be obtained by using "express lines": the activation (or inhibition) of high-level patterns on the descending stream by low-level ascending patterns. Such express lines could activate, for example, "face" models based on partial evidence for a face in the image, and thereby 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 the pathways, in the case of express lines the ascending stream can 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.

This notion of fast initial selection of subsets of models is compatible with a substantial body of psychophysical evidence regarding visual classification and recognition. The first stage in visual recognition appears to be a rough classification of the object into a general class or a small number of classes such as a face, a car, and so on (Rosch et al., 1976). Identification follows in a subsequent stage, and there is evidence (Cavanagh, 1991) that at this stage information associated with the class of objects is used (e.g., to separate actual object contours from cast and attached shadow contours, a distinction that cannot be performed in a bottom-up manner).

The express lines provide one mechanism for "indexing" into the large number of models stored in memory. "Indexing" is a term used in computational vision for the initial selection of a general class, or classes of models, that are likely to correspond to the input image. It appears that such a stage is necessary to enable artificial recognition systems, which are currently limited to a small number of object-models, to deal with large object libraries (Grimson, 1990). The express lines play a role in this process by the fast selection, based on partial information, of likely models on the descending stream. This initial selection is not limited to the activation of models at a single "topmost" level; models at different levels along the descending stream can also be activated and serve as the starting point for descending subsequences. For example, in addition to the selection of a complete face model, intermediate models of face-parts can also be activated. Anatomically, such express lines may correspond to direct connections from low to high visual areas (such as the connections from area V4 to AIT, or from V3 and VP to area TF; Felleman and Van Essen, 1991).

A second mechanism for model selection is provided by the effects of expectation and context. The essential idea is that temporal and spatial correlations can influence the likelihood of different models. Knowledge about the current situation can thereby be used to influence 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 below.

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, for example, under the bed, they may be immediately recognized as a pair of slippers.

Familiar objects can often be recognized in the lack of context, but in dealing with less familiar objects, or with complex scenes, or when the viewing conditions are degraded, the role of context increases in importance and can become indispensable. Even when context is not strictly required, context still facilitates the recognition process, and makes it faster and more reliable (Palmer, 1975; Potter, 1975; Biederman et al., 1982). Context information that helps the observer expect a certain class of objects facilitates recognition significantly, and when objects are placed in an unusual context, recognition is hampered. Under natural conditions, useful context information is almost always present, and this accounts in part for our capacity to deal effectively with complex scenes.

Context effects can operate in the framework of the sequence-seeking scheme by the prior priming of some of the nodes (populations of neurons). The effect will be similar to the mutual priming of the ascending and descending streams but over longer time scales. (Priming between the streams may last for tens to hundreds of milliseconds; context effect should last for considerably longer, up to minutes or hours.) Sequences passing through the primed patterns 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 patterns are activated, for example, by noticing and identifying the bed in the image, they will initiate sequences of their own, and an entire set of patterns will end up in a primed state. Later on, other sequences passing through a primed trace will be facilitated, compared with the nonprimed sequences. The resulting effect is that a context pattern *A* may help to bring about the activation of *B* not as a result of direct prewired association, but because an intermediate subsequence leading from *A* to *B* had been previously facilitated. Context effects will therefore have indirect and widespread influence.

The spread of context effects may 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. It seems that broad, indirect, context effects can be reproduced by the sequence-seeking computation.

Learning Sequences

A simple and local learning rule is sufficient in the counter-streams structure to reinforce selectively complete successful sequences. The reason is that every pattern along a successful sequence will receive both a direct activation and a priming signal from the complementary track; patterns on dead-end tracks will receive one or the other but not both. The approximate temporal coincidence of the two signals can therefore be used to strengthen the successful sequence preferentially. This role is local, since it depends on the activation of a single pattern. Yet, it is sufficient to reinforce preferentially successful sequences forming an uninterrupted link between source and target patterns. Following practice, out of the huge number of possible sequences, those that proved useful in the past will be explored with higher priority in future uses of the network.

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 are indeed involved in the modifications that take place during the learning process (Karni and Sagi, 1991).

In addition to the learning of complete sequences, as above, the system may also be engaged in the learning of the individual stages, that is, the different steps comprising the processing sequences (Poggio, 1990). This aspect of the learning is, however, outside the scope of the present discussion, since the focus is not on the specifics of individual processes, but on their overall common structure.

Due to the parallel exploration of multiple alternatives, and to the tuning of the system by past experience, straightforward recognition tasks will require little or no search. When search becomes necessary, simulations described next suggest that the system will be able to locate the optimal solution efficiently.

Refining the Expansion

The matching between patterns is not an all-or-nothing event, but a graded one: some sequences will lead to better matches than others, and then serve as starting points for exploring additional sequences, that will lead in turn to an improved match. This process has some features in common with a family of optimization and search procedures known as "genetic algorithms" (Holland, 1975), and it is also related to the method of Bayesian optimization (Mockus, 1989). Recent evaluations have shown such methods to behave quite efficiently (Brady, 1985; Peterson, 1990). Our own simulations in the context of pattern matching have also shown that computations based on sequence seeking compare favorably with alternative methods, such as gradient descent and simulated annealing.

Figure 2 shows a simulation of a simplified sequence-seeking process, intended not as a biological model, but to illustrate the process in a simple example. The task is to recognize an input shape (example in Fig. 2*a*) by comparing it with

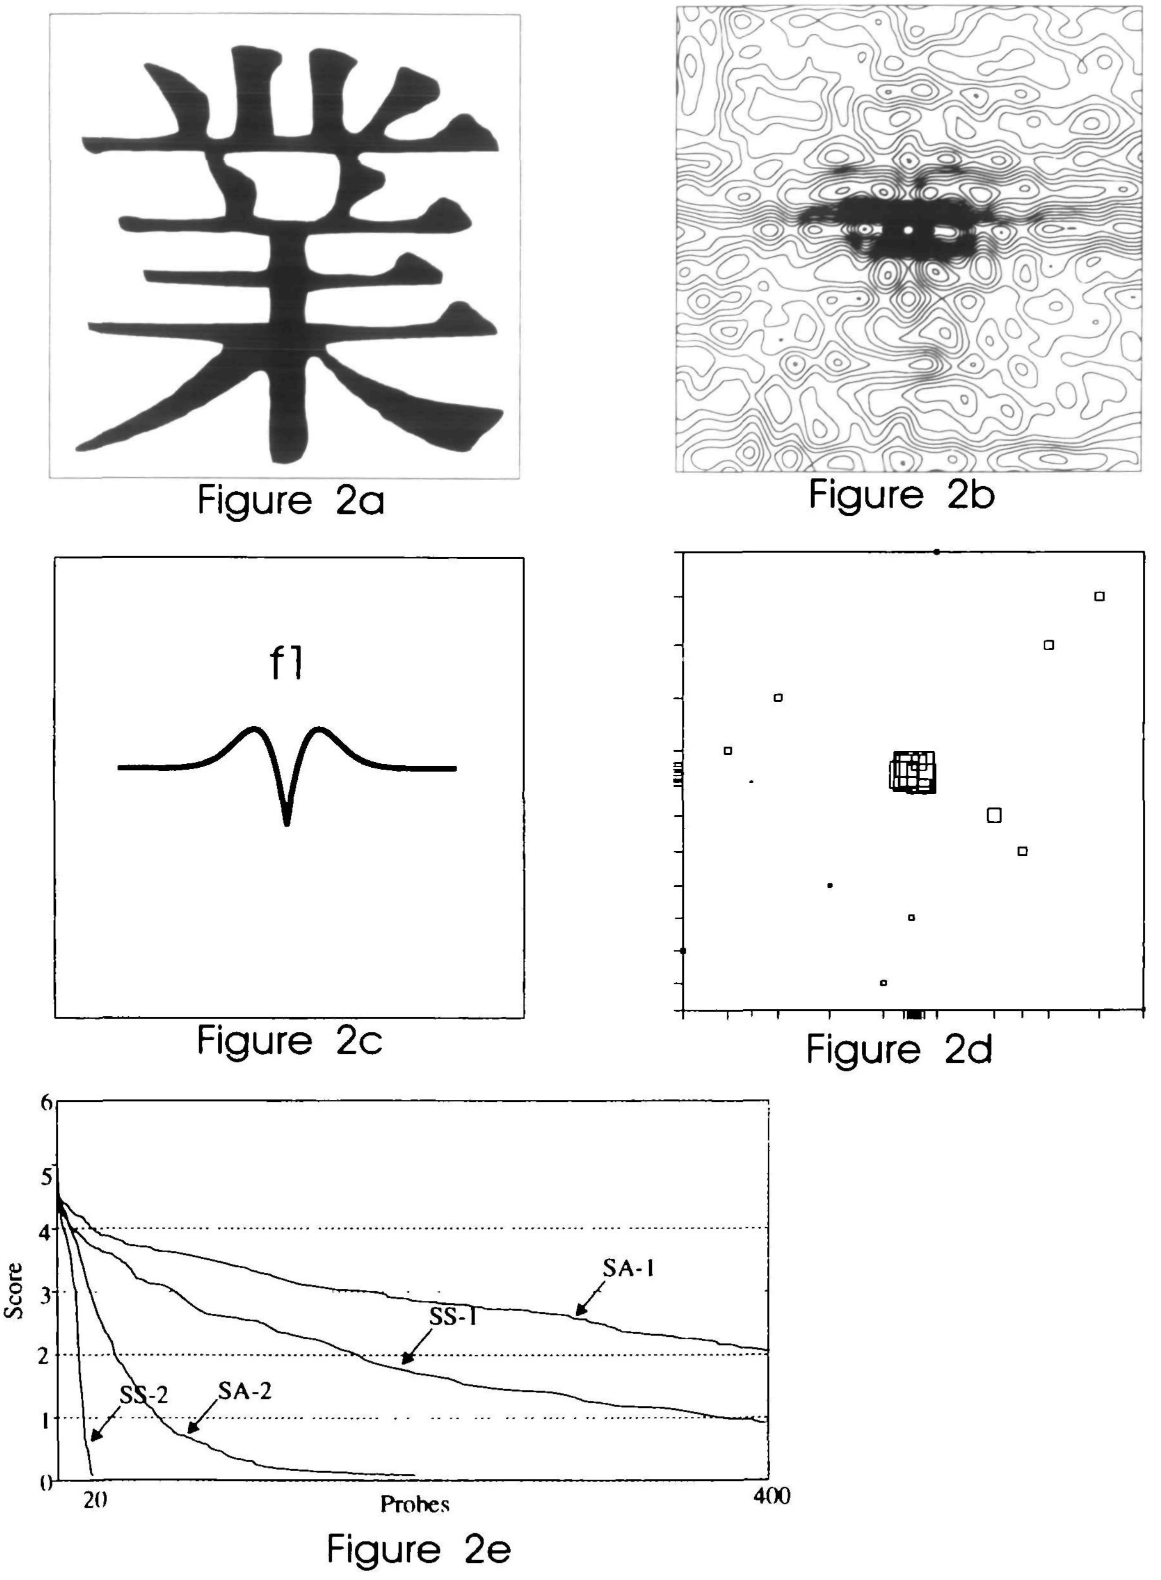
Figure 2a

Figure 2b

Figure 2c

Figure 2d

Figure 2e

Figure 2. Simulation of a simplified sequence-seeking process, used to match an input shape (example in *a*) by comparing it with stored shapes buried in noise. The input shape is displaced in the x,y-directions with respect to the stored pattern. The objective function to be minimized is complex and contains multiple minima (level contours in *b*). In the bidirectional search, the model M is shifted in x, to generate a number of copies M_i, the input image in y, generating displaced copies I_k. A good match between a pair M_i, I_k, leads to the generation of new copies around the corresponding displacement x_i, y_k. In the unidirectional version, M remains fixed, $I_{j,k}$ are generated by shifting I in both x- and y-directions. The likelihood of generating a new sample increases near good past samples, and decreases with the density of past samples. Bias toward good sample points was determined by f_1 (*f1* in *c*), showing the likelihood of generating a new sample point in the vicinity of a previous one. The final likelihood at a point was determined by $\bar{f}_1$, the average of f_1 induced by nearby sample points ($k = 10$) weighted by distance, divided by f_2, a penalty for density, $f_2 = \exp^{1}/_{k} \Sigma^{k}_{i=1} (|x - x_i| + 0.25)^{-2}$. Details of the mathematical derivation of f_1, algorithm, and computational results will be described elsewhere. *d*, A representation of the search following 20 displacements in x and y. The *tick marks* show the selected displacements; the *squares* show the degree of match obtained at some x, y displacements, coded by size. (For each x_i, the best matching y_k is shown, and similarly for the y_js). The search combines exploration of the domain with concentration around good matches. *e*, Match quality as a function of number of patterns explored for one- (SS-1) and two-

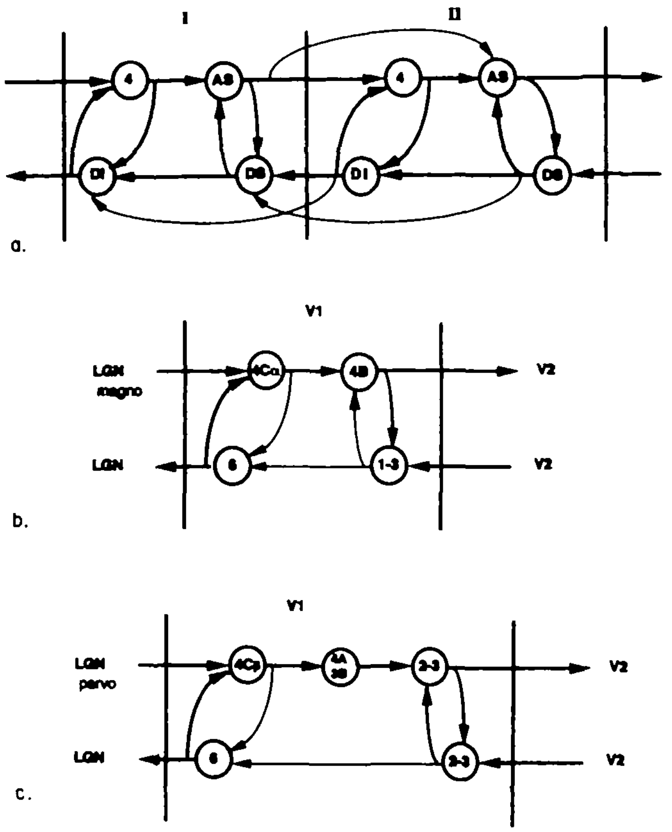

Figure 3. A, How the basic counter-stream structure may be embodied in cortical connectivity. The structure contains two interconnected streams, an ascending and a descending one. The ascending path goes through layer 4 and the ascending superficial population (*AS*) to the next area. The descending path goes from the descending superficial (*DS*) population to *DI* (descending infra) and back to the first area. *Thin arrows* show pathways that "leap over" a step in the stream. Inhibitory and long-range intra-areal connections are not shown. See text for more details. *b*, A schematic representation of the main connections according to the model along the magno stream from the LGN via V1 to V2. (V1 is also connected to other visual areas; not shown in the diagram.) The connections are drawn in a manner suggested by the model and *a*. *Thick arrows*, established connections; *thin arrows*, connections predicted by the model. *c*, The main connections according to the model along the parvo stream from the LGN via V1 to V2. *Thick arrows*, established connections; *thin arrows*, connections predicted by the model.

stored shapes buried in noise. The input shape is displaced in the x, y-directions with respect to the stored pattern. The objective function to be minimized (degree of match as a function of displacement) is complex and contains multiple minima (level contours shown in Fig. 2*b*). A typical search space was composed of about a million possible locations.

To determine the optimal match between the input shape and the stored pattern, one possibility is to displace the input shape by different amounts in x and y, and compare each displaced version with the model, until the best match is obtained. Instead, the simulation used a bidirectional search: the model M was shifted in x, to generate a number of copies M_i, the input image in y, generating displaced copies I_k. The search proceeds by comparing the displaced versions of the input and stored patterns. A good match between a pair M_i, I_k then leads to the generation of new copies around the corresponding displacement x_i, y_k. The results were compared with the unidirectional version, where M remained fixed and $I_{j,k}$ were generated by shifting I in both x- and y-directions.

The procedure used to generate the new "offsprings" around the existing patterns was a simple genetic-like algorithm: the likelihood of generating a new sample increases near good past samples (as determined by the excitatory function f_1 in Fig. 2*c*), and decreases with the density of past samples. This function is computed in the vicinity of past samples. For example, if a good match was obtained between the input displaced horizontally by $\bar{x}$ units and the stored pattern displaced vertically by $\bar{y}$ units, new displacements will be generated around the successful values $\bar{x}$, $\bar{y}$. However, if the match was poor, or if many patterns with similar displacement were already compared, then the likelihood of attempting further solutions in the vicinity of $\bar{x}$, $\bar{y}$ will decrease. In this manner the past samples induce over the search space a likelihood function, and the next samples occur at maxima of this function. The process is simple: it proceeds by trying a number of alternatives, and then selecting and refining successful solutions.

The simulations of the bidirectional search show that the process has a number of favorable general properties. First, the search locates the optimal match efficiently (Fig. 2*e*). It proved more efficient in the pattern matching task than commonly used methods such as simulated annealing (SA in Fig. 2*e*) or gradient descent using multiple starting points. Second, the bidirectional scheme in these examples is considerably more efficient in terms of the number of patterns explored than a unidirectional process (SS-1 vs SS-2 in Fig. 2*e*). This advantage will hold as long as the number of stored patterns to be explored is not too large. Third, the use of past results in guiding the search biases the process to concentrate in more promising regions, compared with SA and gradient descent. Finally, it exhibits good capacity to escape local minima in reaching for the global solution. The example used a simplified task, but the search space was of significant size and contained multiple local minima. Initial experiments using more complex pattern transformations (such as rotation, scaling, and shear) show similar general properties, and additional simulations and mathematical analysis are currently under way.

Generic Aspects of Sequence Seeking

The discussion of the sequence-seeking process used as an example the domain of visual recognition. However, the process of establishing a sequence of transformations, mappings, or states, linking source and target representations, could provide a useful general mechanism for various aspects of perception as well as for nonperceptual 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. Movement trajectories could be based in a sequence-seeking scheme on a repertoire of elementary movements, and these basic movements will then be transformed (scaled, stretched, rotated, etc.) and concatenated together to generate more complex movements. In analogy with sequence seeking in vision, movement planning could also utilize the application of transformations and the generation of compound sequences. Similarly, more general planning and problem solving can also often be formulated in terms of establishing a sequence of

---

←

directional (SS-2) versions of the sequence-seeking search, compared with one- (SA-1) and two-directional (SA-2) versions of simulated annealing. (The score is in units of σ, the standard deviation of the terrain in *b*.)

transformations, mappings, or intermediate states, linking some source and target representations [e.g., Newell and Simon's (1972) GPS model, Quillian's (1968) semantic net theory; see also Winston, 1992], and they may therefore benefit from computations similar to the sequence-seeking scheme. I will not discuss these general problems further, beyond raising the possibility that general aspects of the sequence-seeking process provide a useful computational scheme that could be applied, with appropriate modifications, to different cognitive tasks. This possibility is consonant with the hypothesis, raised by a number of neuroscientists (e.g., Barlow, 1985; Creutzfeldt, 1978; Crick and Asanuma, 1986; Martin, 1988a) regarding the possible existence of some general cortical mechanisms that are applicable, with suitable local modifications, to a broad range of different tasks.

Biological Embodiment

In this section, biological aspects of the model will be discussed. The discussion will focus primarily on general features of the model, such as the ascending and descending streams, and the laminar distribution and general pattern of connections between and within cortical areas. Several aspects of a more specific nature will also be considered, to illustrate possible predictions and questions for further study.

The sequence-seeking model requires two pathways going in opposite directions with the appropriate cross connections. A schematic diagram proposing how the counter-streams structure may be embedded in cortical connections is shown in Figure 3*a*. The proposed embodiment is presented in schematic outline only, focusing on a number of central aspects, but without discussing details or possible variations of the model.

The ascending stream goes through layer 4 to a subpopulation of the superficial layers, denoted in the figure as AS (ascending superficial), and then projects to layer 4 of the next cortical area (II in Fig. 3). The descending stream goes through a different subpopulation of the superficial layers (DS, descending superficial) to DI (descending infra), a subpopulation of the infragranular layers (often in layer 6), and from there to DS of a preceding area. The connections can also leap over one step (or occasionally more) in the stream, for example AS directly to AS on the ascending stream, and DS → DS or 6 → 6 on the descending stream (thin lines in Fig. 3*a*).

Layer 5 is left out of the diagram because, according to the model, it (or a part of it) is involved primarily not in the main streams, but with their control, in cooperation with subcortical structures. There are at least two reasons for assuming that layer 5 (or parts of it, e.g., layer 5b of the macaque's V1) may be involved in control functions. First, its orderly connections to subcortical structures (e.g., from visual cortex to the pulvinar and the superior colliculus, structures implicated in controlling attention and eye movements; Desimone et al., 1990) that are reciprocally connected in turn in a topographic manner to multiple visual areas. Second, the firing pattern of a population of pyramidal cells in this layer that "can initiate synchronized rhythms and project them on neurons in all layers" (Silva et al., 1991, p 434).

Note that the counter-streams structure suggests a natural organization in about five or six main layers: one or two performing control functions, two (an input and an output layer) for the ascending, and two for the descending streams. The division between the roles of the different layers is likely to be in reality less clear cut, and there are known variations and specialized sublaminations. However, the main goal of the diagram is to emphasize the possible common underlying structure according to the model, rather than to account for possible variations.

It is interesting to note that from a developmental standpoint the layered cortical structure appears to develop in two stages, possibly from distinct origins (Marin-Padilla, 1978; Deacon, 1990). The most superficial and deepest layers develop first, and all other layers develop subsequently in between them. These two structures may be the precursors of the descending and ascending streams, respectively (Deacon, 1990; Mumford, 1994). This developmental view is compatible with the notion of the two distinct and interconnected streams of the counter-streams structure.

Connections of V1: Data and Predictions

To give a more specific example, Figure 3, *b* and *c*, shows an expanded version of the diagram, applied to cortical area V1 (which is somewhat special, but for which the data are more comprehensive than for other visual areas), and its connections to the LGN and cortical area V2 (V1 is also connected to other visual areas, not shown in the diagram). Figure 3, *b* and *c*, shows the connections in the macaque of the magnocellular stream and the parvocellular stream, respectively (Rockland and Lund, 1983; Lund, 1988a,b; Martin, 1988a). The diagram shows the main connections; additional secondary ones will not be considered. The connections are drawn in a manner suggested by the model, and they include both known connections (thick arrows) and connections predicted by the proposed scheme but for which empirical evidence is partial or lacking (thin arrows).

The pattern of connections in the two streams appears to be in general agreement with the counter-streams structure and Figure 3*a*. The model suggests that a similar structure can be used, with local modifications, as a repeating circuit within a large network to utilize the cortex inherent parallelism and combine ascending and descending information flows.

If the general hypothesis regarding the counter-streams structure is broadly correct, then a number of predictions can be made regarding the main connectivity patterns within and between areas. One general prediction is the possible distinction between the AS and DS subpopulations. This separation reflects the most straightforward implementation of the scheme; however, some alternatives can exist without violating the constraints of the model.

A separation between the 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. In the magnocellular projection from V1 to V2 the forward projection originated mainly in layer 4B, while the back-projection is mainly to other layers (Fig. 3*b*). It is further expected that even when the superficial layers provide both the source and the target of connections to another area, there will in fact often be a separation to the AS/DS subpopulations. If these populations exist, they should be connected in a reciprocal manner. A related expectation derived from the model is the existence of priming-type synaptic interactions, that is, excitatory synaptic input that by itself may not be very effective in driving the target cells, but that facilitates the effects of subsequent inputs to these cells.

An example at the other end of the spectrum, of a specific expectation, is that in the magnocellular stream the model suggests 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), and that the same superficial cells connected to layer 4B will also be the recipients of descending projections from V2 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.

The model also includes a reciprocal connection between layer 4, and the LGN-projecting cells in layer 6. The projection from layer 6 to layer 4 is well-established in both the cat (McGuire et al., 1984) and monkey (Lund, 1988a), and there is support for the opposite connection as well (Lund and Boothe, 1975). It is 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 et al., 1984), in accordance with the connectivity in Figure 3, *b* and *c*.

The connections between layers 4 and 6 are expected in the model to have a priming effect (not necessarily the only effect; see Bolz and Gilbert, 1986; Martin, 1988b), and this notion has some physiological support. It was found (Ferster and Lindström, 1985) that using electrical activation of layer 6 cells by antidromic activation increased the probability of layer 4 firing, and most cells fire multiple spikes in response to each stimulus. Under the opposite conditions, when layer 6 was inactivated, the main observed effect was the reduction in excitability of layer 4 cells (Grieve et al , 1991).

From an anatomical standpoint, EM reconstructions (McGuire et al., 1984) have shown terminations of layer 6 axons on smooth and sparsely spiny cells. The model suggests also a projection onto layer 4 spiny cells. Recent data by Ahmed et al. (1994) support this suggestion, and indicate that the major target of layer 6 pyramidal output to layer 4 is in fact the spiny stellate cells.

Layer 6 is also involved in the model in the descending pathway (although, as mentioned, layers 5 and 6 are often further divided into distinct sublayers, and the identification of layer 6 with the DI population in the model is not always straightforward). The involvement of layer 6 in the descending pathway, as either the origin or the target of the descending projection, has been demonstrated in various visual areas, including V1, V2, V3, V4, TEO, MT, MST, VIP, PO, and LIP (Lund et al., 1981; Maunsell and Van Essen, 1983; Felleman and Van Essen, 1984; Rockland and Virga, 1989; Andersen et al., 1990; Rockland et al., 1992).

Lateral Connections between Areas

Connections between cortical areas (not only visual, 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 (Rockland and Pandya, 1979; Friedman, 1983; Maunsell and Van Essen, 1983; Van Essen, 1985; Zeki and Shipp, 1988; Andersen et al., 1990; Boussaud et al., 1990; Felleman and Van Essen, 1991). Lateral connections terminate in all layers, and their origin is bilaminar, from the supra as well as infra layers. The lateral pattern is relatively complex; it is therefore interesting that a number of its main features can be derived almost directly from the model. The counter-streams structure does not have a distinct, third type of connections. It allows, however, forward and backward connection simultaneously in both directions, and it can include lateral connections by assuming that they are the union of ascending and descending connections. If this view is correct, then the main connections participating in the lateral connection can be inferred from the basic scheme (Fig. 3*a*). According to the model, they include the direct connections, AS → 4 and DI → DS, as well as the connections that leap over one stage in the diagram, AS → AS, DS → DS,DI, and DI → DI.

The origin of the projections according to the model would be bilaminar, and the terminations would span all layers, in agreement with the observed pattern. This can also explain several difficulties such as the problem of irregular terminations (Felleman and Van Essen, 1991) that occurs, for example, when some of the terminations are restricted to layer 4 of the target area while others show columnar terminations. This was termed F/C (i.e., a mixture of "four" and "columnar") paradoxical termination, since termination in layer 4 is a signature for ascending connections, while a columnar termination signifies lateral connections. In the counter-streams structure, the point to note is that the lateral connections from the superficial layers of area *A* to target area *B* are composed of two subprojections: AS → 4 (ascending) and DS → DS,DI (descending). Anterograde labeling of the upper layers of area *A* can therefore show mixed patterns of terminations, such as layer 4 alone, or a columnar termination, in agreement with the F/C paradoxical termination. It can also (by labeling the DS alone) show a bilaminar pattern of connections; this can account for the other types of irregular terminations.

The detailed nature of these connections is still not entirely conclusive. However, the proposed account serves to illustrate two points; first, that some of the apparent complexities may have a natural explanation within the counter-stream structure; second, if the account is generally correct, it provides support for the existence of the AS and DS subpopulations in the model.

Priming Mechanisms

Synaptic interactions in the model include priming-type effects between the complementary streams. Although this has not been studied directly, some known or physiologically plausible mechanisms could play a role in such priming interactions.

Priming can be obtained for example by long-lasting depolarization, combined with subsequent input, added either linearly or nonlinearly. A long-lasting depolarization can be caused by a number of possible mechanisms, including ionic channels with a slow time course (McCormick, 1990; Hirsch and Gilbert, 1991; Amitai et al., 1993), NMDA receptors (Miller et al., 1989), or the activation of distal parts of the dendritic tree (Stratford et al., 1989). This depolarization will facilitate subsequent inputs by summation (Miller et al., 1989), or by a nonlinear interaction (Koch, 1987; Esguerra et al., 1989; Sherman et al., 1990). Although the details are not known, it appears that synaptic mechanisms for priming connections are physiologically plausible, and it will be of interest to try to test them empirically.

Effects of the Feedback Projection

According to the sequence-seeking scheme, the physiological effects of the descending projections can assume two different forms: either the priming and modulation of the ascending stream, or the direct activation of a lower area.

Both effects have been observed in physiological studies, modulatory (Sandell and Schiller, 1982; Nault et al., 1990), as well as direct excitatory effects (Cauller and Kullics, 1991; Mignard and Malpeli, 1991). Further predictions of the model regarding the modulatory effects include (1) that similar modulatory effects are also likely to be induced by ascending signals on descending ones, and (2) that the two effects may be segregated into two distinct subpopulations: in Figure 1*c*, $\bar{B}$ can be directly driven along the descending stream, but patterns such as *B* on the ascending stream are expected to show modulatory effects.

In summary, the computation proposed by the sequence-seeking model is a bidirectional process performed by top-down and bottom-up streams of processing seeking to meet. Bottom-up processing is supported by the ascending pathways; top-down processing, by the descending ones. In each direction, different alternatives are explored in parallel. The scheme incorporates a number of basic lessons from computer vision and perceptual psychology. Essential properties of the scheme include the simultaneous exploration of multiple alternatives, the relatively simple, uniform, and extensible structure, the flexible use of "bottom-up" and "top-down" sequences that can meet at any level, the roles of context and of fast classification, and the learning of complete sequences by a simple local reinforcement rule.

The model combines the proposed computation with a number of known as well as predicted aspects of cortical circuitry. Given the still limited knowledge regarding cortical structures and the computations they perform, the model addresses mainly general aspects of the computation. The combination of the proposed computation and structure serves to suggest a framework, that offers a computational account for several basic features of cortical circuitry, such as the predominantly reciprocal connectivity between cortical areas, the forward, backward, and lateral connection types, the regularities in the distribution patterns of interarea connections, the organization in five or six layers, and the effects of backprojections, as well as a number of more specific details. It also poses problems for further study at the structural as well as computational levels.

Notes

This work was supported in part by NSF Grant IRI-8900267 and in part by the Israel Science Foundation administered by the Israel Academy of Sciences and Humanities. I thank C. Gilbert, E. Grimson, E. Hildreth, C. Koch, J. Maunsell, D. Mumford, T. Poggio, and W. Richards for helpful comments.

Correspondence should be addressed to Shimon Ullman, Department of Applied Mathematics and Computer Science, The Weizmann Institute of Science, P.O. Box 266, Rehovot 76100, Israel.

References

Abeles M (1991) Corticonics. Cambridge: Cambridge UP.

Ahmed B, Andersen JC, Douglas RJ, Martin KAC, Nelson JC (1994) Polyneural innervation of spiny stellate neurons in cat visual cortex. J Comp Neurol 339:1-11.

Amitai Y, Friedman A, Connors BW, Gutnick MJ (1993) Regenerative activity in apical dendrites of pyramidal cells in neocortex. Cereb Cortex 3:26-38.

Andersen RA, Asanuma C, Essick G, Siegel RM (1990) Cortico-cortical connections of anatomically and physiologially defined subdivisions within the inferior parietal lobule. J Comp Neurol 296:65-113.

Barlow HB (1985) Cerebral cortex as model builder. In: Models of the visual cortex (Rose D, Dobson VG, eds), pp 37-46 New York: Wiley.

Biederman I, Mazzanotte RJ, Rabinowitz JC (1982) Scene perception: detecting and judging objects undergoing relational violations. Cognit Psychol 14:143-147.

Bolz J, Gilbert CD (1986) Generation of end-inhibition in the visual cortex via interlaminar connections. Nature 320:362-365.

Boussaud D, Ungerleider LC, 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 RM (1985) Optimization processes gleaned from biological evolution. Nature 317:804-806.

Cauller LJ, Kullics AT (1991) The neural basis of the behaviorally relevant N1 component of the somatosensory-evoked potential in SI cortex of the awake monkey: evidence that backward cortical projections signal conscious touch sensation. Exp Brain Res 84. 607-619.

Cavanagh P (1991) What's up in top-down processing? In: Representations of vision (Gorea A, ed), pp 295-304. Cambridge: Cambridge UP.

Creutzfeldt OD (1978) The neocortical link: thoughts on the generality of structure and function of the neocortex. In: Architectonics of the cerebral cortex (Brazier MAB, Petsche M, eds), pp 357-380. New York: Raven

Crick F, Asanuma C (1986) Certain aspect of the anatomy and physiology of the cerebral cortex. In: Parallel distributed processing, Vol 2 (McClelland JL, Rumelhart DE, eds), pp 333-371. Cambridge, MA: MIT Press.

Deacon TW (1990) Rethinking mammalian brain evolution. Am Zool 30:629-706.

Desimone R, Wessinger M, Thomas L, Schneider W (1990) Attentional control of visual processing: cortical and subcortical mechanisms. Cold Spring Harbor Symp Quant Biol 55:963-971.

Distler C, Ungherleider LG, Boussaud D, Desimone R (1991) Cortical connections of temporal-lobe area TEO in macaque. Eur J Neurosci [Suppl] 4:56

Edelman GM (1978) Group selection and phasic re-entrant signalling: a theory of higher brain functions. In: The mindful brain (Edelman GM, Mountcastle VB, eds), pp 51-100. Cambridge, MA: MIT Press.

Engel AK, König P, Kreiter AK, Schillen TB, Singer W (1992) Temporal coding in the visual cortex: new vistas on integration in the nervous system. Trends Neurosci 15:218-226.

Esguerra M, Kwon YH, Sur M (1989) NMDA and non-NMDA receptors mediate retinogeniculate transmission in cat and ferret LGN *in vitro*. Soc Neurosci Abstr 15:175

Felleman DJ, Van Essen DC (1984) Cortical connections of area V3 in macaque extrastriate cortex. Soc Neurosci Abstr 10:933

Felleman DJ, Van Essen DC (1991) Distributed hierarchical processing in primate visual cortex. Cereb Cortex 1:1-47.

Ferster D, Lindström S (1985) Synaptic excitation of neurones in area 17 of the cat by intracortical axon collaterals of cortico-geniculate cells. J Physiol (Lond) 367.233-252.

Freuder EC (1974) AI technical report 345, A computer vision system for visual recognition using active knowledge. Cambridge, MA: MIT

Friedman DP (1983) Laminar patterns of termination of corticocortical afferents in the somatosensory system. Brain Res 273:147-151.

Fukushima K (1986) A neural network model for selective attention in visual pattern recognition. Biol Cyber 55:5-15.

Geman S, Geman D (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. IEEE Transact Pattern Anal Machine Intell 6:721-741.

Grieve KL, Murphy PC, Sillito AM (1991) Inhibitory and excitatory components to the subcortical and cortical influence of layer VI cells in the cat visual cortex. Soc Neurosci Abstr 17:629.

Grimson WEL (1990) Object recognition by computer. Cambridge, MA: MIT Press.

Hirsch JA, Gilbert CD (1991) Synaptic physiology of horizontal connections in the cat's visual cortex. J Neurosci 11:1800-1809.

Holland JH (1975) Adaptation in natural and artificial systems. Ann Arbor: University of Michigan.

Jordan M, Jacobs RA (1993) MIT computational cognitive science technical report 9301, Hierarchical mixtures of experts and the EM algorithm. Cambridge, MA: MIT.

Karni A, Sagi D (1991) Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. Proc Natl Acad Sci USA 88:4966-4970.

Katz LC, Burkhalter A, Dreyer WJ (1984) Fluorescent latex microspheres as a retrograde neuronal marker for *in vivo* and *in vitro* studies of visual cortex. Nature 310:498-500.

Koch C (1987) The action of the corticofugal pathway on sensory thalamic nuclei: a hypothesis. Neuroscience 23:399-406.

Lowe DG (1985) Perceptual organization and visual recognition. Boston: Kluwer.

Lund JS, Hendrickson AE, Orgen MP, Tobin EA (1981) Anatomical organization of primate visual cortex VII. J Comp Neurol 202.19-45.

Lund JS (1988a) Excitatory and inhibitory circuitry and laminar mapping strategies in primary visual cortex of the monkey. In: Signal and sense: local and global order in perceptual maps (Edelman GM, Gall WE, Cowan WM, eds), pp 51-66. New York: Wiley.

Lund JS (1988b) Anatomical organization of macaque monkey striate visual cortex. Annu Rev Neurosci 11:253-288.

Lund JS, Boothe RG (1975) Interlaminar connections and pyramidal neuron organization in the visual cortex, area 17, of the macaque monkey *Macaca mulatta*. J Comp Neurol 159:305-334.

Marin-Padilla M (1978) Dual origin of the mammalian neocortex and evolution of the cortical plate. Anat Embryol (Berl) 152:109-126.

Marr D (1982) Vision. San Francisco: Freeman.

Martin KAC (1988a) From single cells to simple circuits in the cerebral cortex. Q J Exp Physiol 73:637-702.

Martin KAC (1988b) The lateral geniculate nucleus strikes back. Trends Neurosci 11:192-194.

Matan O, Burges CJC, Le Cun Y, Denker JS (1992) Multi-digit recognition using a space displacement neural network. In: Neural Information Processing Systems, 4 (Moody J, Hanson S, Lippmann R, eds), pp 488-495. San Mateo, CA: Kaufmann

Maunsell JHR, Gibson JR (1992) Visual response latencies in striate cortex of the macaque monkey. J Neurophysiol 68:1332-1344.

Maunsell JHR, Van Essen DC (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.

McGuire BA, Hornung J-P, Gilbert CD (1984) Patterns of synaptic input to layer 4 of cat striate cortex J Neurosci 4:3021-3033.

Mignard M, Malpeli JG (1991) Paths of information flow through visual cortex. Science 251:1249-1251.

Miller KD, Chapman B, Stryker MP (1989) Visual responses in adult cat visual cortex depend on *N*-methyl-D-aspartate receptors. Proc Natl Acad Sci USA 86:5183-5187.

Mockus J (1989) Bayesian approach to global optimization. Boston: Kluwer.

Mumford D (1992) On the computational architecture of the neocortex II. The role of cortico-cortical loops. Biol Cybern 66:241-251.

Mumford D (1994) Neural architectures for pattern-theoretic problems. In: Large-scale neuronal theories of the brain (Koch C, Davis JL, eds), in press. Cambridge, MA: MIT Press.

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

Newell A, Simon HA (1972) Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.

Okajima K (1991) A recurrent system incorporating characteristics of the visual system: a model for the function of backward neural connections in the visual system. Biol Cybern 65:235-241.

Palmer SE (1975) The effects of contextual scenes on the identification of objects. Mem Cognit 3:519-526.

Peterson C (1990) Parallel distributed approaches to combinatorial optimization: benchmark studies on traveling salesman problem. Neural Comput 2:261-269.

Poggio T (1990) A theory of how the brain might work. In: The Brain: CSH Symposia on Quantitative Biology, pp 899-910. New York: CSH Laboratory.

Poggio T, Torre V, Koch C (1985) Computational vision and regularization theory. Nature 317:314-319.

Potter MC (1975) Meaning in visual search. Science 187:565-566.

Quillian MR (1968) Semantic memory. In: Semantic information processing (Minsky M, ed), pp 227-270. Cambridge, MA: MIT Press.

Rockel AJ, Hiorns RW, Powell TPS (1980) The basic uniformity in structure of the neocortex. Brain 103:221-244.

Rockland KS, Lund JS (1983) Intrinsic laminar lattice connections in primate visual cortex. J Comp Neurol 216:303-318.

Rockland KS, Pandya DN (1979) Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. Brain Res 179:3-20.

Rockland KS, Virga A (1989) Terminal arbors of individual "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.

Rockland KS, Saleem K, Tanaka K (1992) Widespread feedback connections from areas v4 and TEO. Soc Neurosci Abstr 18:390.

Rolls ET, Tovee MJ, Lee B (1991) Temporal response properties of neurons in the macaque inferior temporal cortex. Eur J Neurosci [Suppl] 4:84.

Rosch E, Mervis CB, Gray WD, Johnson DM, Boyes-Braem P (1976) Basic objects in natural categories. Cognit Psychol 8:382-439.

Sandell JH, Schiller PH (1982) Effect of cooling area 18 on striate cortex cells in the squirrel monkey. J Neurophysiol 48:38-48.

Sejnowski TJ (1986) Open questions about computation in cerebral cortex. In: Parallel distributed processing, Vol 2 (McClelland JL, Rumelhart DE, eds), pp 372-389. Cambridge, MA: MIT Press.

Sherman SM, Scharfman HE, Lu SM, Guide W, Adams PR (1990) *N*-methyl-D-aspartate (NMDA) and non-NMDA receptors participate in EPSPs of cat lateral geniculate neurons recorded in thalamic slices. Soc Neurosci Abstr 16:159.

Silva LR, Amitai Y, Connors BW (1991) Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. Science 251:432-435.

Sporns O, Tononi G, Edelman GM (1991) Modeling perceptual grouping and figure-ground segregation by means of active reentrat connections Proc Natl Acad Sci USA 88:129-133.

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

Tenenbaum JM, Barrow HG (1976) Technical report 123, Experiments in interpretation-guided segmentation. Stanford, CA: Stanford Research Institute.

Thorpe SJ, Celebrinin S, Trotter Y, Imbert M (1991) Dynamics of stereo processing in area V1 of the awake primate. Eur J Neurosci [Suppl] 4:83.

Tononi G, Sporns O, Edelman GM (1992) Reentry and the problem of integrating multiple cortical areas: simulation of dynamic integration in the visual system. Cereb Cortex 2:310-335.

Ullman S (1989) Aligning pictorial descriptions: an approach to object recognition. Cognition 32:193-254.

Ullman S (1991) MIT AI memo 1331, Sequence seeking and counter streams: a model for information processing in the cortex. Cambridge, MA: MIT.

Ullman (1993) The visual recognition of three-dimensional objects. In: Attention and performance XIV (Meyer DE, Kornblum S, eds), pp 79-98. Cambridge, MA: MIT Press.

Van Essen DC (1985) Functional organization of primate visual cortex. In: Cerebral cortex, Vol 3 (Peters A, Jones EGJ, eds), pp 259-329. New York: Plenum.

White EL (1989) Cortical circuits. Boston: Birkhäuser.

Winston PH (1992) Artificial intelligence, 3d ed. Reading, MA: Edison Wesley.

Yuille A, Hallinan P (1992) Deformable templates. In: Active vision (Blake A, Yuille A, eds), pp 21-38. Cambridge, MA: MIT Press.

Zeki S, Shipp S (1988) The functional logic of cortical connections. Nature 335:311-317.

Zipser D, Rumelhart DE (1990) The neurobiological significance of the new learning models. In Computational neuroscience (Schwartz EL, ed), pp 192-200. Cambridge, MA: MIT Press.