

HOW DOES THE CEREBRAL CORTEX WORK?

Learning, Attention, and Grouping by the Laminar Circuits of Visual Cortex

Stephen Grossberg*

Department of Cognitive and Neural Systems
and
Center for Adaptive Systems
Boston University
677 Beacon Street
Boston, MA 02215

April, 1998

Spatial Vision, in press

To appear in the special issue on Long-Range Visual Mechanisms

Technical Report CAS/CNS-97-023

Running Title: Cortical Circuits for Learning, Attention, and Grouping

* Supported in part by the Defense Advanced Research Projects Agency, the National Science Foundation (NSF IRI-97-20333), and the Office of Naval Research (ONR N00014-95-1-0409 and ONR N00014-95-1-0657).

S. Grossberg

Acknowledgments: The author wishes to thank Robin Amos and Diana Meyers for their valuable assistance in the preparation of the manuscript.

The organization of neocortex into layers is one of its most salient anatomical features. These layers include circuits that form functional columns in cortical maps. A major unsolved problem concerns how bottom-up, top-down, and horizontal interactions are organized within cortical layers to generate adaptive behaviors. This article models how these interactions help visual cortex to realize: (1) the binding process whereby cortex groups distributed data into coherent object representations; (2) the attentional process whereby cortex selectively processes important events; and (3) the developmental and learning processes whereby cortex shapes its circuits to match environmental constraints. New computational ideas about feedback systems suggest how neocortex develops and learns in a stable way, and why top-down attention requires converging bottom-up inputs to fully activate cortical cells, whereas perceptual groupings do not.

Introduction. The cerebral cortex is the seat of our highest forms of intelligence, and its understanding is thus a goal for all students of mind and brain. Neocortex has an intricate design which exhibits a characteristic organization into six distinct cortical layers (Brodmann, 1909; Martin, 1989). Differences in the thickness of these layers and the sizes and shapes of neurons led the German anatomist Korbinian Brodmann to identify more than fifty divisions, or areas, of neocortex. This classification has been invaluable to later scientists as a basis for discerning different functional roles for different parts of the brain. On the other hand, why the neocortex has six layers, or indeed a laminar design, has remained a mystery from a functional point of view. The present article proposes a model that provides clear functional roles for these layers for purposes of visual perception, and suggests that similar functional roles may be at work in all sensory and cognitive processing.

Linking cortical anatomy to behavior cannot be done without a sufficiently powerful method. This is true because cortical organization exhibits multiple scales of processing, including individual neurons within the various layers, neural circuits that link these neurons within and between these layers, functional columns that are defined through these interlaminar interactions, cortical maps that are defined by the global organization of these columns within a cortical area, and thalamocortical and corticocortical interactions that occur between different thalamic and cortical areas. These cortical interactions, moreover, occur both bottom-up, from more peripheral to more central areas, and top-down, from more central to more peripheral areas, and have a characteristic laminar organization of their own.

In order to make functional sense of such complex interactions, one needs to be able to link cortical organization to the behaviors that it controls; one needs to show how these designs lead to useful behavioral properties that have been selected and maintained through evolution. The present article describes such a linkage. It shows that perceptual, attentional, developmental, and learning properties that are known to be carried out by visual cortex place severe, and seemingly contradictory, demands on cortical organization. When one tries to realize all of these constraints within a single, unified cortical architecture, a model of how this is accomplished can be identified, whose every component is supported by neurobiological data. In addition to providing a unified explanation of many behavioral and neurobiological data for which no alternative explanation has yet arisen, this cortical model also makes a number of testable predictions. The constraints that will be described herein are all well-known. They have not been simultaneously applied before in part because they are usually studied experimentally by different groups of scientists, and also because their remarkably constraining nature does not become apparent until one attempts to model all of them at once.

Perceptual Grouping and Attention: A Paradox. Some of the main constraints can be stated in terms of perceptual processes that are familiar to us all. During visual perception, the visual cortex can generate perceptual groupings and can focus attention upon objects of interest. *Perceptual grouping* is the process whereby the brain organizes image contrasts into emergent boundary structures that segregate objects and their backgrounds in response to texture, shading and depth cues in scenes and images (Julesz, 1971; Ramachandran and Nelson, 1976; Beck, Prazdny, and Rosenfeld, 1983; Polat and Sagi, 1994). Perceptual grouping is a basic step in solving the “binding problem”, whereby spatially distributed features are bound into representations of objects and events in the world. Illusory contours are a particularly vivid form of perceptual grouping. Figure 1 shows how an illusory contour can form over image positions that do not receive contrastive bottom-up inputs from an image or scene. Perceptual groupings can form *preattentively* and automatically, without requiring the conscious attention of a viewing subject.

Attention enables humans and other animals to selectively process information that is of interest to them. In contrast to perceptual grouping, top-down attention does not form visible percepts over positions that receive no bottom-up inputs. Attention can sensitize, or prime, an observer to expect an object to occur at a given location, or with particular stimulus properties (Posner, 1980; Duncan, 1984). But attention, by itself, is no substitute for the object’s actual appearance. Were attention able to routinely generate fully formed perceptual representations, we could not tell the difference between external reality and internal fantasy.

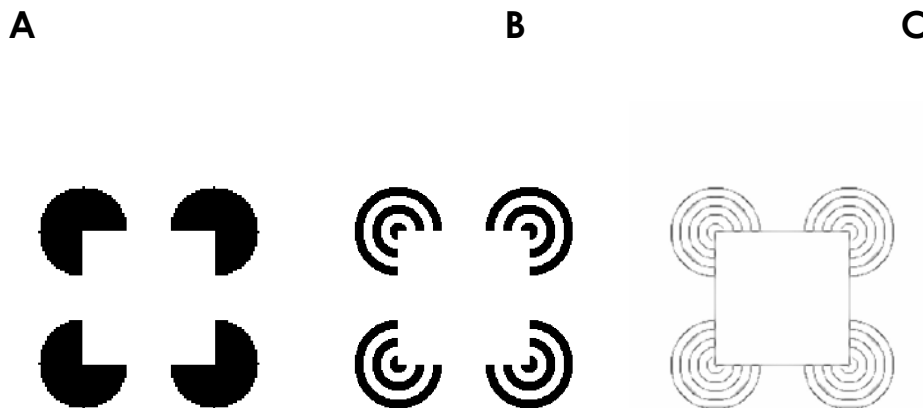


Figure 1. An illusory Kaniza square can be perceived (A) colinear to edge inducers and (B) perpendicular to line-end inducers. (C) Model simulation of the latter type of boundary grouping (reprinted with permission from Grossberg, Mingolla, and Ross, 1997).

Given that perceptual grouping and attention make opposite requirements on bottom-up inputs, the question of how they are simultaneously realized within the same circuits of the visual cortex is a difficult one to answer. One possible answer to this question is that these circuits are *not* simultaneously realized within the same cortical areas. This answer is not, however, supported by recent cortical data. For example, it has been shown that short-range perceptual groupings can occur within cortical area V1 (Redies, Crook, and Creutzfeldt, 1986; Grosf, Shapley, and Hawken, 1993) and that longer-range perceptual groupings can occur within cortical area V2

(Von der Heydt, Peterhans, and Baumgartner, 1984; Peterhans and von der Heydt, 1989). In addition, attentional focussing occurs from the earliest visual cortical area V1 top-down to the Lateral Geniculate Nucleus, or LGN (Sillito, Jones, Herstein, and West, 1994). Attention also operates in visual cortical areas V1, V2, and V4 (Motter, 1994a, 1994b; Reynolds, Nicholas, Chelazzi, and Desimone, 1995; Beauchamp, Cox, and DeYoe, 1997; Hupé *et al.*, 1997; Ito, Westheimer, and Gilbert, 1997; Johnson and Burkhalter, 1997; Lamme, Zipser, and Spekreijse, 1997; McAdams and Maunsell, 1997; Press and van Essen, 1997) and areas MT and MST (O'Craven *et al.*, 1997; Treue and Maunsell, 1997). Many recent neurophysiological experiments show that attentional processing operates throughout the visual cortex.

Both perceptual grouping and attentional modulation are thus integrated within the visual cortical areas V1 and V2. How does this circuitry form perceptual groupings that *can* complete a boundary grouping over locations which receive no bottom-up visual inputs, whereas top-down attention *cannot* do so? Why *should* attention be deployed throughout the visual cortex, including cortical areas which previously were thought to accomplish purely preattentive processing? An answer can be found by exploring the link between attention and learning, and using this link to further constrain the model.

The Link Between Attention and Learning. Earlier modeling work has suggested that top-down attention is a key mechanism whereby the brain solves the *stability-plasticity* dilemma (Grossberg, 1980, 1982; Grossberg and Stone, 1986; Carpenter and Grossberg, 1991, 1993; Grossberg, 1995; Grossberg and Merrill, 1996). The stability-plasticity dilemma concerns that fact that our brains can rapidly learn enormous amounts of information throughout life, without just as rapidly forgetting what they already know. Brains are *plastic* and can rapidly learn new experiences, without losing the *stability* that prevents catastrophic forgetting.

The fact that multiple thalamic and cortical levels develop through experience-dependent learning (Hubel, Wiesel, and LeVay, 1977; Stryker and Harris, 1986; Calloway and Katz, 1990; Antonini and Stryker, 1993a, 1993b; DeAngelis, Ohzawa, and Freeman, 1993; Ghose, Freeman, and Ohzawa, 1994; Galuske and Singer, 1996) raises the question of how such attentive processes may be realized within neocortex in order to stabilize this learning through time. This question, in turn, leads to further constraints on cortical design, because at least some perceptual groupings can form preattentively, and provide the substrate upon which higher-level attentional processes can act. How can the preattentive grouping mechanisms develop in a stable way, before any higher-order attentional processes can develop with which to stabilize them? Why does not this problem lead to an infinite regress; namely, why is it not the case that attentional mechanisms cannot develop until preattentive mechanisms do, but preattentive mechanisms cannot develop stably in the absence of attention? I call this the *attention-preattention interface problem*. It is an *interface* problem because it is shown below how laminar cortical circuits enable preattentive grouping processes to use some of the same circuitry that attentive mechanisms use, even before attentive mechanisms may come into play, in order to stabilize their own cortical development and learning.

The solution proposed herein to the attention-preattention interface problem builds upon earlier efforts to solve the stability-plasticity dilemma. Adaptive Resonance Theory, or ART, proposes a solution of how attention solves the stability-plasticity dilemma by modeling how bottom-up signals activate top-down expectations whose signals are matched against bottom-up data (Figure 2A). Both the bottom-up and top-down pathways contain adaptive weights, or long-term

memory traces, that may be modified by experience. The learned top-down expectations “focus attention” upon information that matches them (Figure 2B). They select, synchronize, and amplify the activities of cells within the attentional focus, while suppressing the activities of irrelevant cells, which could otherwise be incorporated into previously learned memories and thereby destabilize them.

The cell activities which survive such top-down attentional focusing rapidly reactivate bottom-up pathways, thereby generating a type of feedback *resonance* between bottom-up and top-down signal exchanges (Figure 2A). Such resonances rapidly bind distributed information at multiple levels of brain processing into context-sensitive representations of objects and events. These resonances are proposed to support slower processes of learning; hence the name *adaptive resonance*.

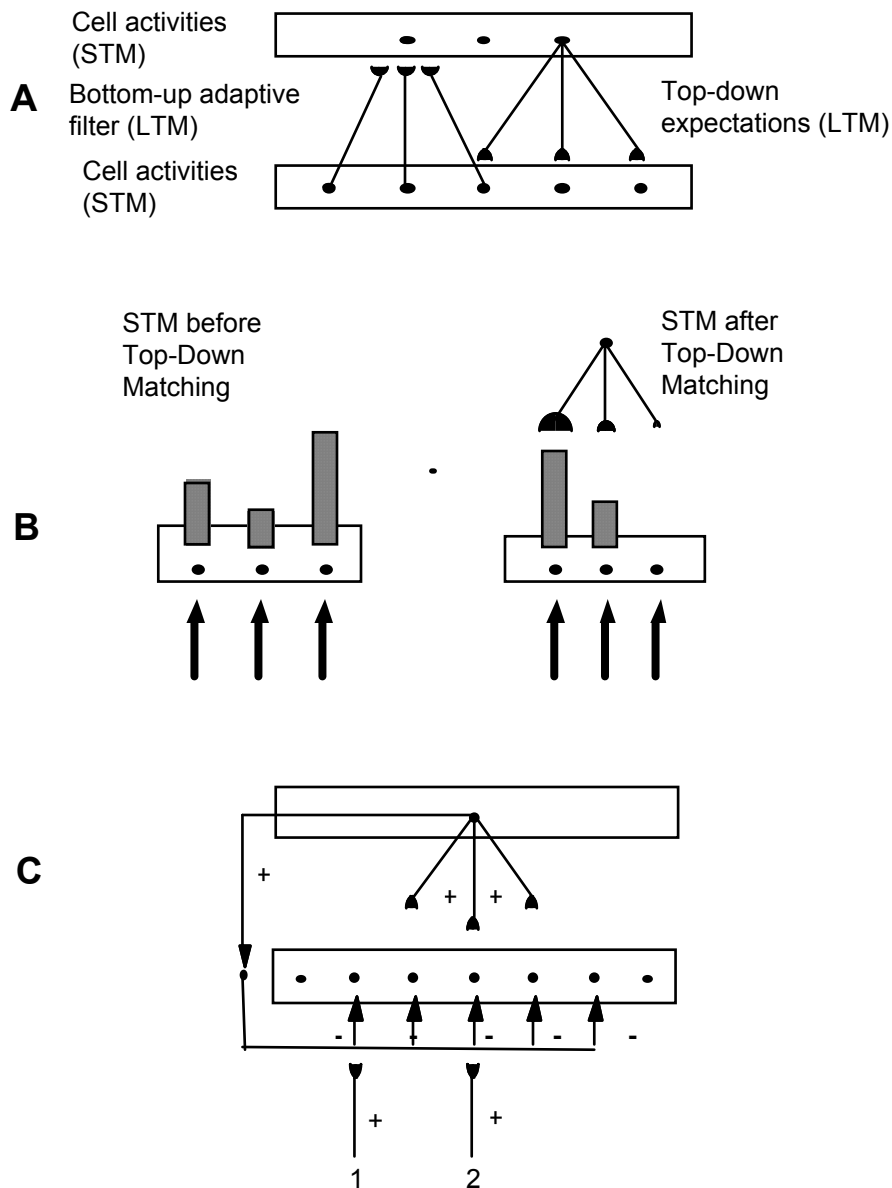


Figure 2. (A) Patterns of activation, or short-term memory (STM), on a lower processing level send bottom-up signals to a higher processing level. These signals are multiplied by adaptive weights, or learned long-term memory (LTM) traces, which influence the activation of the cells at the higher processing level. These latter cells, in turn, activate top-down expectation signals that are also multiplied by learned LTM traces. These top-down expectations are matched against the STM pattern that is active at the lower level. (B) This matching process confirms and amplifies STM activations that are supported by large LTM traces in an active top-down expectation, and suppresses STM activations that do not get top-down support. The size of the hemidisks at the end of the top-down pathways represents the strength of the learned LTM trace that is stored in that pathway. (C) The ART Matching Rule may be realized by a top-down on-center off-surround network, as discussed in the text.

ART has shown, using mathematical proofs and data explanations, how the learning of receptive field properties during early development, and the learning of perceptual and cognitive representations during adulthood, could easily suffer catastrophic forgetting in response to a changing world (Grossberg, 1976a, 1976b; Carpenter and Grossberg, 1987), and how top-down attention can stabilize learning if it satisfies four properties (Carpenter and Grossberg, 1987, 1991), which together are called the ART Matching Rule:

Bottom-Up Automatic Activation: A cell, or cell population, can become active enough to generate output signals if it receives a large enough bottom-up input, other things being equal. Such an input can drive the cell to supraliminal levels of activation.

Top-Down Priming: A cell becomes subliminally active if it receives only a large top-down expectation input. Such a top-down priming signal can sensitize, or modulate, the cell, and thereby prepare it to react more quickly and vigorously to subsequent bottom-up inputs that match the top-down prime. The top-down prime by itself cannot, however, generate supraliminal output signals from the cell.

Match: A cell becomes active if it receives large convergent bottom-up and top-down inputs. Such a matching process can generate enhanced activation as resonance takes hold.

Mismatch: A cell's activity is suppressed, even if it receives a large bottom-up input, if it also receives only a small, or zero, top-down expectation input.

ART originally proposed four related circuits that could all realize these properties (Carpenter and Grossberg, 1987) and thereby stabilize the learning process. More recent data analyses have suggested that variants of the simplest circuit, a top-down on-center off-surround network (Figure 2C), is used by the brain (Grossberg, 1995). Figure 2C clarifies how such a circuit can achieve all four properties. In particular, when only bottom-up signals are active, all cells can fire that receive large enough inputs. When only top-down attention is active, cells that receive inhibition but no excitation get inhibited, while cells that receive a combination of excitation and inhibition can get at most subliminally activated due to the balance between excitation and inhibition. When bottom-up and top-down inputs match, as in pathway 2 of Figure 2C, the two excitatory sources of excitation that converge at the cell can overwhelm the one inhibitory source; it is a case of "two-against-one." When bottom-up and top-down inputs mismatch, as in pathway 1 of Figure 2C, the top-down inhibition can neutralize the bottom-up excitation; it is a case of "one-against-one."

Is Attention Too Slow and Weak to Matter? Many scientists have resisted the concept that top-down attention plays an important role, despite the fact that there are massive top-down pathways throughout the cortex (Macchi and Rinvik, 1976; Tsumoto, Creutzfeldt, and Legendy, 1978; van Essen and Maunsell, 1983; Felleman and van Essen, 1991), due to the belief that feedback cannot operate quickly enough to influence cortical information processing. Two mathematical properties of ART systems show such concerns to be groundless: First, cortical interactions that obey ART properties can choose the correct cells on the first pass of bottom-up signalling, if the input pattern is familiar to the system; second, resonance can stabilize within even a single processing cycle of attentional feedback (Carpenter and Grossberg, 1987, 1991; Grossberg and Somers, 1991).

Another roadblock to understanding derives from the properties, explained by ART, that top-down attention accomplishes subliminal, or subthreshold, *priming* and *matching*. By itself, it cannot supraliminally activate cells, thereby enabling them to generate output signals, and thus seems to be too “weak” to significantly affect cortical processing. For example, Zeki and Shipp (Zeki and Shipp, 1988, p. 316) wrote that “backward connections seem not to excite cells in lower areas, but instead influence the way they respond to stimuli”. Likewise, the data of Sillito *et al.* (1994, pp. 479-482) on attentional feedback from V1 to LGN led them to conclude that “the cortico-thalamic input is only strong enough to exert an effect on those dLGN cells that are additionally polarized by their retinal input...the feedback circuit searches for correlations that support the ‘hypothesis’ represented by a particular pattern of cortical activity”. Their experiments demonstrated all of the properties of the ART Matching Rule, since they found in addition that “cortically induced correlation of relay cell activity produces coherent firing in those groups of relay cells with receptive-field alignments appropriate to signal the particular orientation of the moving contour to the cortex...this increases the gain of the input for feature-linked events detected by the cortex”. In other words, top-down priming, by itself, cannot fully activate LGN cells; it needs matched bottom-up retinal inputs to do so; and those LGN cells whose bottom-up signals support cortical activity get synchronized and amplified by this feedback. In addition, anatomical studies have shown that the top-down V1 to LGN pathway realizes a top-down on-center off-surround network (Dubin and Cleland, 1977; Weber, Kalil, and Behan, 1989), as in Figure 2.

What Stabilizes Perceptual Development and Learning? The above discussion suggests that top-down attentional mechanisms should be present in *every* cortical area wherein learning can occur, since without top-down learned expectations that focus attention via the ART Matching Rule, any such learned memories could easily be degraded due to catastrophic forgetting.

These analyses should, in particular, apply to the perceptual grouping process, because the cortical horizontal connections that support perceptual grouping in areas like V1 develop through a learning process that is influenced by visual experience (Luhmann, Martinez Millan, and Singer, 1986; Calloway and Katz, 1990; Gilbert and Wiesel, 1992; Lowel and Singer, 1992; Antonini and Stryker, 1993a; Galuske and Singer, 1996). It is also known that many developmental and learning processes, including those that control horizontal connections, are stabilized dynamically, and can be reactivated by lesions and other sources of cortical imbalance (Das and Gilbert, 1995; Gilbert and Wiesel, 1992), and that adult learning uses the same types of mechanisms as the infant developmental processes upon which it builds (Bailey *et al.*, 1992; Kandel and O’Dell, 1992; Mayford *et al.*, 1992). What cortical mechanisms ensure this type of dynamical stability?

This is a particularly challenging problem for the case of perceptual groupings for two reasons which, on the surface, seem to be unrelated, but which will turn out to be intimately related on a mechanistic level. The first reason is that perceptual groupings can form *preattentively*. How, then, can attention control their stability during infant development and adult learning? This is the *attention-preattention interface problem* that was mentioned above. The second reason, also noted above, is that perceptual groupings can form over positions that do not receive bottom-up inputs, as in the case of illusory contours. They therefore seem to violate the ART Matching Rule. How, then, can the horizontal connections that generate perceptual groupings maintain themselves in a stable way? Why are they not washed away whenever a grouping forms over positions that do not receive a bottom-up input?

My proposed answer to this question unifies two types of neural models which have been developed along separate paths for two decades: the attentive ART model, and the preattentive perceptual grouping model that is called the Boundary Contour System, or BCS (Grossberg and Mingolla, 1985; Grossberg, 1994; Gove, Grossberg, and Mingolla, 1995). Until the present, it has not been possible to understand how ART circuits are embodied within the laminar architecture of visual cortex. Recent work on the BCS has suggested how preattentive grouping may be carried out by laminar visual cortex. My proposed synthesis of how attention, perceptual grouping, development, and perceptual learning are realized by the laminar circuits of visual cortex builds upon this new foundation. The present article focuses on one combination of intracortical and intercortical pathways. It does not attempt to model all cortical connections, or the variations that exist across species. Preliminary studies indicate, however, that the principles modeled herein can be elaborated and adapted to handle these cases.

Analog Coherence of Perceptual Groupings. This BCS model proposes how the laminar circuitry of visual cortex enables perceptual groupings to maintain their analog and spatial context-sensitivity in response to changes in stimulus properties (Grossberg, Mingolla, and Ross, 1997). Analog sensitivity means that perceptual groupings can alter their form in response to graded changes in stimulus properties. For example, one boundary grouping, such as an illusory contour, may form if some signals are weak and others strong, whereas a different grouping may form if the reverse relative signals strengths occur. Spatial context-sensitivity means that the cortex can respond flexibly to spatial rearrangements of the stimuli that are to be grouped.

The grouping process actively selects and binds together the most salient groupings for conscious perception, while suppressing less salient groupings. This selection, or binding, process endows each grouping with an inner coherence, so that object representations are not merely the sum of their features. Unfortunately, processes that select winning groupings while inhibiting losing groupings tend to wipe out analog sensitivity; they tend to generate binary outcomes. Because analog values carry useful information about objects and events, it is important to understand how the cortex can bind distributed information into coherent representations without a loss of analog sensitivity. Remarkably, the laminar circuitry of visual cortex is capable of robustly realizing the key property of *analog coherence* whereby groupings form coherently without losing analog or spatial context-sensitivity.

Four Preattentive Designs of Visual Cortex. Four circuit properties summarize this proposal of how the visual cortex, notably areas V1 and V2, uses its laminar design to generate perceptual groupings that preserve analog coherence. Each design principle will be described along with cortical data that it explains. Then four more circuit properties will be proposed whereby attention, development, and learning are integrated into this laminar design.

Design 1. Analog Sensitivity to Bottom-Up Sensory Inputs. Bottom-up inputs from the retina go through the Lateral Geniculate Nucleus (LGN) on their way to cortex. LGN outputs directly excite layer 4 (Hubel and Wiesel, 1962; Chapman, Zaks, and Stryker, 1991; Reid and Alonso, 1995). LGN inputs also excite layer 6, which then indirectly influences layer 4 via an on-center off-surround network of cells (Ferster and Lindstrom, 1985; Grieve and Sillito, 1991a, 1991b, 1995), as in Figure 3A. The net effect of LGN inputs on layer 4 cells is thus via an on-center off-surround network. Such a feedforward on-center off-surround network of cells can preserve the analog sensitivity of, and normalize, the activities of target cells if these cells obey the membrane equations of neurophysiology (Grossberg, 1973, 1980; Douglas *et al.*, 1995). In the present case, such a network preserves the analog sensitivity of layer 4 cells in response to LGN inputs that may vary greatly in intensity.

Design 2. Bipole Boundary Grouping. The active layer 4 cells input to pyramidal cells in layer 2/3. These cells initiate the formation of perceptual groupings. They generate excitatory signals among themselves using monosynaptic long-range horizontal connections, and inhibition using short-range disynaptic inhibitory connections (Hirsch and Gilbert, 1991; McGuire *et al.*, 1991), as in Figure 3B. These interactions support inward perceptual groupings between two or more boundary inducers (von der Heydt, Peterhans, and Baumgartner, 1984; Peterhans and von der Heydt, 1989), but not outward groupings from a single inducer (Hirsch and Gilbert, 1991; Cannon and Fullenkamp, 1993; Knierim and van Essen, 1992; Somers, Nelson, and Sur, 1995; Stemmler, Usher, and Niebur, 1995), as in the case of illusory contours (Figure 1). If a single

S. Grossberg

inducer could generate groupings, our percepts would become crowded with webs of boundaries that spread out from every feature in a scene.




Figure 3. A model circuit of retinal, lateral geniculate nucleus (LGN), and cortical V1 interactions: Open symbols indicate excitatory interactions and closed symbols inhibitory interactions. (A) Feedforward circuit from retina to LGN to cortical layers 4 and 6: Retina: Retinal ON cells have an on-center off-surround organization. Retinal OFF cells have an off-center on-surround organization. LGN: The LGN ON and OFF cells receive feedforward ON and OFF cell inputs from the retina. Layer 4: Layer 4 cells receive feedforward inputs from LGN and layer 6. LGN ON and OFF cell excitatory inputs to layer 4 directly establish oriented simple cell receptive fields. Layer 6 cells excite layer 4 cells with a narrow on-center and inhibit them using inhibitory interneurons that span a broader off-surround. Like-oriented layer 4 simple cells with opposite contrast polarities compete (not shown) before generating half-wave rectified outputs that converge on layer 2/3 pyramidal (complex) cells. Layer 2/3: The converging simple

cell outputs enable complex cells to respond to both polarities. They hereby full-wave rectify the image. [Figure caption continued on next page.]

Figure 3 Continued. (B) Horizontal grouping interactions in layer 2/3: After being activated by inputs from layer 4, layer 2/3 pyramidal (complex) cells excite each other monosynaptically via horizontal connections, primarily on their apical dendrites. They also inhibit one another via disynaptic inhibition that is mediated by model smooth stellate cells. Multiple horizontal connections share a common pool of stellate cells near each target pyramidal cell. This ensures that boundaries form inwardly between pairs or greater numbers of boundary inducers, but not outwardly from a single inducer. (C) Cortical feedback loop from Layer 2/3 to Layer 6: Layer 6 cells receive excitatory inputs from layer 2/3. The long-range cooperation hereby engages the feedforward layer 6-to-4 on-center off-surround network, which then reactivates layer 2/3 cells. This “folded feedback” loop can select winning groupings without a loss of analog coherence. (D) Outputs from layer 2/3 to area V2 directly excite layer 4 cells and layer 6 cells, which indirectly influence layer 4 cells via an on-center off-surround network, as in area V1.

These grouping properties are ensured in the following way: When a single active pyramidal cell sends horizontal monosynaptic excitation to other pyramidal cells, this excitation is inhibited by the disynaptic inhibition that it also generates; this is another case of “one-against-one”. A different result obtains when two or more pyramidal cells are activated at positions that are located at opposite sides of a target pyramidal cell, and all the cells are approximately colinear across space. Then the excitation from the active pyramidal cells summates at the target cell, thereby generating a larger total excitatory input than a single pyramidal cell could. In addition, the active cells all excite a single population of disynaptic inhibitory interneurons, which generates a saturating, or normalized, inhibitory output to the target cell. Thus excitation is bigger than inhibition in this case, so that grouping can occur; it is another case of “two-against-one.” This combination of constraints is called the *bipole* property. Layer 2/3 pyramidal cells may hereby become active either due to direct inputs from layer 4, or due to bipole boundary groupings that form in response to other active layer 2/3 cells.

Design 3. Folded Feedback and Analog Coherence. The active cells in layer 2/3 send excitatory feedback signals to layer 6 via layer 5 (Gilbert and Wiesel, 1979; Ferster and Lindstrom, 1983), as in Figure 3C. Layer 6, in turn, once again activates the on-center off-surround network from layer 6 to 4. This feedback process is called *folded feedback*, because feedback signals from layer 2/3 to layer 6 get transmitted in a feedforward fashion back to layer 4. The feedback is hereby “folded” back into the feedforward flow of bottom-up information within the laminar cortical circuits.

Folded feedback turns the cortex into a feedback network that binds the cells throughout layers 2/3, 4, and 6 into functional columns (Mountcastle, 1957; Hubel and Wiesel, 1962, 1977). The on-center off-surround network helps to select the strongest groupings that are formed in layer 2/3 and to inhibit weaker groupings, while preserving the analog values of the selected groupings. In particular, the on-center signals from layer 6-to-4 support the activities of those pyramidal cells in layer 2/3 that are part of the strongest horizontal groupings. The off-surround signals inhibit inputs to layer 4 that were supporting less active groupings in layer 2/3. In this way, signals from layer 4 to the less active groupings in layer 2/3 are removed, and thus these groupings collapse.

Design 4. Self-Similar Hierarchical Boundary Processing. Converging evidence suggests that area V2 replicates the structure of area V1, but at a larger spatial scale (Kisvarday *et al.*, 1995). In particular, layer 2/3 in area V1 sends bottom-up inputs to layers 4 and 6 of area V2, much as LGN sends bottom-up inputs to layers 4 and 6 of area V1 (van Essen and Maunsell, 1983; Felleman and van Essen, 1991); see Figure 3D. This input pattern from V1 to V2 can preserve the analog sensitivity of layer 4 cells in V2 for the same reason that the LGN inputs to V1 can preserve the analog sensitivity of layer 4 cells in V1. The shorter perceptual groupings in layer 2/3 of area V1 (Redies, Crook, and Creutzfeldt, 1986; Grosz, Shapley, and Hawken, 1993) are proposed to group together, and thereby enhance the signal-to-noise ratio of, nearby V1 cells with similar orientation and disparity selectivity. The longer perceptual groupings in area V2 (van der Heydt, Peterhans, and Baumgartner, 1984; Peterhans and von der Heydt, 1989) are proposed to build long-range boundary segmentations that separate figure-from-background; generate 3-D groupings of the edges, textures, shading, and stereo information that go into object representations; and complete boundaries across gaps in bottom-up signals due to the retinal blind spot and veins (Grossberg, 1994; Grossberg and McLoughlin, 1997). Both types of groupings achieve analog coherence by using the same type of laminar circuitry. Computer simulations of such groupings are reported elsewhere (Grossberg, Mingolla, and Ross, 1997).




Figure 4. (A) Top-down corticogeniculate feedback from Layer 6: LGN ON and OFF cells receive topographic excitatory feedback from layer 6 in V1, and more broadly distributed inhibitory feedback via LGN inhibitory interneurons that are excited by layer 6 signals. The feedback signals pool outputs over all cortical orientations and are delivered equally to ON and OFF cells. Corticogeniculate feedback selects, gain-controls, and synchronizes LGN cells that are consistent with the cortical activation that they cause, thereby acting like a type of automatic attentional focus. (B) Attentional feedback from V2 to V1: Layer 6 in V2 activates layer 6 in V1, which then activates the layer 6-to-4 on-center off-surround network that attentionally primes layer 4 cells. (C) One feedback pathway arises from Layer 6 cells in V2 and activates apical dendrites in Layer 1 of V1. Cells in Layer 5 are activated through these apical dendrites and thereupon activate Layer 6 cells.

Four Attentive Designs of Visual Cortex. How does top-down attention fit into these layered circuits without disturbing the preattentive grouping process? Four attentive design principles are

proposed to accomplish this, and to enable cortical development and learning to proceed in a stable fashion.

Design 5. Top-Down Feedback from V1 to LGN. As noted above, layer 6 of area V1 sends a top-down on-center off-surround network to the LGN (Murphy and Sillito, 1987; Weber, Kalil, and Behan, 1989; Murphy and Sillito, 1996), as in Figure 4A. This top-down pathway automatically focuses attention on those LGN cells whose activities succeed in activating V1 cells. Data of Sillito *et al.* (1994) show that this feedback obeys the ART Matching Rule, and thus can only subliminally activate, or modulate, LGN cells. Matched bottom-up inputs are needed to supraliminally activate LGN cells while top-down signals are active. This process is predicted to help stabilize the development of disparity-tuned complex cells in V1 during the visual critical period.

Design 6. Folded Feedback from Layer 6 of V2 to Layer 4 of V1. A variant of this top-down process seems to occur at all stages of visual cortex, and probably beyond. Layer 6 in a given cortical area, such as V2, generates top-down cortical signals to layer 6 of lower cortical areas, such as V1, where they activate the layer 6-to-4 folded feedback network in the lower area (Figure 4B). One such known top-down pathway exits layer 6 in V2 and activates V1 via layer 1 (Pandya and Yeterian, 1985). This pathway activates layer 1 apical dendrites of layer 5 cells, which relay them to layer 6 cells in V1 (Cauller and Connors, 1994; Rockland, 1994), as in Figure 4C. Top-down feedback hereby activates a top-down on-center off-surround circuit, much like the ART circuit in Figure 2. I propose that it is realized in cortex using outputs from layer 6 of a given cortical area to activate layer 4 of a lower cortical area via layer 6-to-4 folded feedback. This proposal is supported by neurophysiological data showing that top-down signals activate the center and inhibit the surround of area V1 cells (Bullier *et al.*, 1996).

Design 7. Layer 6-to-4 Signals are Subliminal. The ART Matching Rule predicts that this top-down pathway *subliminally* activates, or modulates, cells in layer 4. I propose that this subliminal property is due to the fact that the on-center off-surround interactions from layer 6-to-4 are balanced so that at most a weak excitatory effect occurs after activating the circuit via top-down feedback. This hypothesis is consistent with neurophysiological data from ferret visual cortex showing that the layer 6-to-4 circuit is functionally weak (Wittmer, Dalva, and Katz, 1997). In addition, Hupé *et al.* (1997, p. 1031) note: “feedback connections from area V2 modulate but do not create center-surround interactions in V1 neurons.” Thus top-down feedback from layer 6 of V2 is predicted to be able to supraliminally activate layer 6 of V1 but not layer 4 of V1.

This functional “weakness” does not prevent the top-down circuit from having a major effect on cortical cell activations when the cortex is activated bottom-up by visual inputs. In particular, by inhibiting activities of layer 4 cells whose layer 2/3 cell projections are not bound into strong groupings, the layer 6-to-4 circuit can dramatically change the *balance* of activation across the cortex, thereby enabling the strongest groupings to become amplified and resonant. This analysis predicts that attentional and grouping constraints from higher levels of cortex can feed back to selectively bias the groupings that arise at lower cortical levels. In particular, the “higher-order” boundary completion and figure-ground perception grouping properties of V2 can select cells in V1 which are consistent with them.

Design 8. Two Bottom-Up Input Sources to Layer 4. A simple functional explanation can now be provided of why there are direct bottom-up inputs to layer 4, as well as indirect bottom-up inputs to layer 4 via layer 6, in many cortical areas (e.g., Figures 3A and 3D). Why are not these two separate input pathways redundant? Why, in particular, is not the indirect layer 6-to-4 pathway sufficient to activate layer 4 cells *and* to maintain their analog sensitivity using its on-center off-surround network? The proposed explanation is that direct inputs to layer 4 are needed to supraliminally activate layer 4 cells, since the indirect layer 6-to-4 inputs, being subliminal, cannot do so.

Taken together, these eight cortical design principles lead to the circuit diagram for perceptual grouping and attention between LGN, V1, and V2 that is shown in Figure 5. I propose that the same cortical circuits may explain data at multiple levels of cortical organization; for example, they are sufficient to explain attentional data collected from macaque cortical areas V2 and V4 (Motter, 1994a, 1994b; Reynolds et al., 1995), wherein top-down attention once again selectively primes features within the attentional focus while suppressing the effects of features that are not.




Figure 5. A model synthesis of bottom-up, top-down, and horizontal interactions in LGN, V1, and V2. Cells and connections in green denote preattentive excitatory mechanisms that are involved in perceptual grouping. Red denotes inhibitory mechanisms. Blue denotes top-down attentional mechanisms.

Stable Cortical Development and Adult Learning. With the circuit of Figure 5 in mind, we can return to the fundamental question of how the horizontal connections within cortical area V1 develop through a learning process that is modifiable by visual inputs. How is such development stabilized through time?

A simple and parsimonious answer can now be given, because both preattentive perceptual groupings within V1 and attentive feedback from V2 to V1 both generate feedback signals to layer 6 of V1. Both types of feedback activate the folded feedback circuit from layer 6-to-4. Top-down attention uses this circuit to focus attention within V1 by inhibiting layer 4 cells that are not supported by excitatory 6-to-4 feedback. Perceptual grouping uses it to inhibit layer 4 cells that would otherwise help to form incorrect groupings. In so doing, folded feedback prevents the wrong combinations of cells in layers 4 and 2/3 from being active simultaneously, and thereby prevents incorrect horizontal connections from being learned between these active cells.

The folded feedback circuit from layer 6-to-4 gets activated by perceptual grouping signals from layer 2/3-to-6 at *all* positions of the grouping, even those positions that do not receive bottom-up inputs. The ART Matching Rule is thus satisfied at all such positions, and the source of the “top-down expectation” is the perceptual grouping itself. In summary, the *preattentive perceptual grouping is its own attentional prime* (Figure 6).




Figure 6. The preattentive perceptual grouping is its own attentional prime: An intracortical perceptual grouping, such as an illusory contour, uses [Figure caption continued on next page.]

Figure 6 Continued. the same layer 6-to-4 on-center off-surround network as does intercortical attentional priming. Its off-surround can hereby inhibit incorrect layer 4 cells that could otherwise enter into incorrect associations with active layer 2/3 cells. This mechanism works at all positions of the grouping, even those that do not receive bottom-up inputs.

Model simulations of how cortical area V1 develops its horizontal connections in layer 2/3 and its interlaminar connections between layers 6 and 4 have demonstrated that stable development can be achieved if the on-center off-surround circuits from layer 6-to-4 subliminally activate layer 4 cells (Grossberg and Williamson, 1997). When this constraint is not realized, it is much harder to stabilize development in the model. By extension, these results suggest how perceptual learning in the adult can be stabilized using the same mechanisms.

Some Predictions. Many predictions are implicit in the cortical design principles that are proposed above, since each one predicts a functional role for a known cortical circuit. Some particularly worthwhile tests would be: (1) test if top-down V1 to LGN feedback plays a role in stabilizing the development of disparity tuning during the visual critical period; (2) test if top-down feedback from V2 to V1 can supraliminally activate layer 6 cells, but only subliminally activate layer 4 cells; (3) test if a long-range horizontal grouping in layer 2/3 of V2 can inhibit vertically oriented simple cells at the midpoint of this grouping in layer 4 of V1; (4) test if layer 4 simple cells cannot be supraliminally activated if only the LGN-to- 6-to-4 input pathway is active. Each of these predictions is technically difficult to test, but as increasingly selective methods are developed (e.g., Nelson et al., 1994), this will become easier to do.

Many additional tests of the model can be envisaged. Another test incorporates habituating or depressing synapses (Abbott, Varela, Sen, and Nelson, 1997; Grossberg, 1969, 1976b, 1980). Several earlier articles have used the hypothesis that habituating synapses may exist in the spatial competition stage that occurs prior to bipole grouping, in order to explain psychophysical data about visual persistence, perceptual aftereffects, and residual traces (Francis and Grossberg, 1996a, 1996b; Francis, Grossberg, and Mingolla, 1994). In the present laminar model of cortex, such synapses may be interpreted to occur in the layer 6-to-4 on-center off-surround network. If this hypothesis is correct, then the following type of experiment, along with variants thereof, may be of interest. First, use direct electrical stimulation of area V2 to generate, say, a horizontally oriented top-down priming signal to layer 6 of area V1. If such a signal generates suprathreshold output signals from layer 6, then they should be able to habituate the synapses in the layer 6-to-4 on-center off-surround network. This habituating effect may then be tested with bottom-up horizontal and vertical visual stimuli to the same location. If the top-down habituating effect works in the same manner as a prior bottom-up adaptation stimulus (Meyer, Lawson, and Cohen, 1975), then persistence of cell responses in layers 2/3 and 4 will be less with a horizontal stimulus and greater with a vertical stimulus than it would have been in the absence of the top-down habituating prime. Such a test would exploit the hypotheses that both the intercortical top-down priming circuit and the bottom-up activation circuit utilize the same layer 6-to-4 on-center off-surround network, and that habituating synapses exist in this network. A related test in area V2 would habituate the layer 6-to-4 horizontally oriented synapses in the middle of an illusory contour and would then do the same horizontal and vertical persistence test as described above. This test would exploit the hypotheses that intracortical feedback signals from layer 2/3 of V2 and feedforward signals from area V1 to V2 utilize the same layer 6-to-4 on-center off-surround

network, and that habituated synapses exist in this network. Taken together, experiments such as these would provide further information about how every preattentive perceptual grouping becomes its own attentional prime.

Learning, Attention, and Grouping in Sensory and Cognitive Neocortex. The present article suggests how bottom-up, top-down, and horizontal interactions are organized within visual cortical areas V1 and V2. The key impediment to making this synthesis in the past derived from an absence of hypotheses about how the cortex uses the folded feedback circuit from layer 6-to-4 to achieve both top-down attentional priming and analog coherence of preattentive perceptual groupings. The hypothesis that this priming circuit obeys the ART Matching Rule enables all previous results about how ART interactions stabilize development and learning to be applied to the case of cortical development and adult perceptual learning. Many data about preattentive grouping of information, attentive priming of information, infant development, and adult learning may now be functionally linked to the laminar, columnar, and map organization of the visual cortex. This is true, for example, of recent neurophysiological data which have supported the key ART predictions that many developmental and learning processes are stabilized dynamically using learned top-down expectations, and that adult learning uses the same types of mechanisms as the infant developmental processes upon which it builds. These data include: shared molecular substrates of neonatal development and adult learning (Bailey *et al.*, 1992; Kandel and O'Dell, 1992; Mayford *et al.*, 1992); attentional modulation of cortical development (Singer, 1982); plasticity of adult cortical representations after lesions (Merzenich *et al.*, 1988); dynamical reorganization of long-range connections in the visual cortex (Gilbert and Wiesel, 1992; Zohary *et al.*, 1994); fast perceptual learning in the adult (Karni and Sagi, 1991; Poggio, Fahle, and Edelman, 1992); and fast cortical synchronization (Eckhorn *et al.*, 1988; Gray and Singer, 1989).

Given the generality of the functional problems that these circuits are proposed to solve, it seems plausible that similar circuits may be used throughout the neocortex. Indeed, ART models have already been used, albeit without a laminar cortical interpretation, to explain developmental, cognitive, and neurobiological data about normal and amnesic recognition learning, categorization, working memory, memory search, and hypothesis testing (Grossberg, 1980, 1982; Grossberg and Stone, 1986; Carpenter and Grossberg, 1991, 1993; Grossberg, 1995; Grossberg and Merrill, 1996; Grossberg, Boardman, and Cohen, 1997). In many of these cognitive examples, an orienting system, which has been proposed to be at least partly realized in the hippocampal system, interacts with the attentional thalamocortical circuits that do the learning. This interaction drives hypothesis testing, or memory search, for new, or more task-appropriate, recognition categories within the attentional system. Such an orienting system enables attention to be allocated in a more flexible way than can be achieved by the attentional mechanisms on their own, and permits an enormous expansion in the amount of information that can be learned. In this regard, long-range intrinsic connections are also known to occur in many areas of neocortex; for example, in the auditory and language areas of the human temporal cortex (Schmidt *et al.*, 1997). It remains to be seen whether these model circuits generalize to show how other sensory and cognitive regions of the neocortex are functionally organized.

References

- Abbott, L.R., Varela, K., Sen, K., and Nelson, S.B. (1997). Synaptic depression and cortical gain control. *Science* **175**, 220-223.
- Antonini, A. and Stryker, M.P. (1993a). Functional mapping of horizontal connections in developing ferret visual cortex: Experiments and modeling. *Journal of Neuroscience* **14**, 7291-7305.
- Antonini, A. and Stryker, M. P. (1993b). Rapid remodeling of axonal arbors in the visual cortex. *Science* **260**, 1819-1821.
- Bailey, C. H., Chen, M., Keller, F., and Kandel, E. R. (1992). Serotonin-mediated endocytosis of a pCAM: An early step of learning-related synaptic growth in aptysia. *Science* **256**, 645-649.
- Beauchamp, M. S., Cox, R. W., and DeYoe, E. A. (1997). Gradients of attention in the human visual motion processing system. *Society for Neuroscience Abstracts* **179.3**, 457.
- Beck, J., Prazdny, K., and Rosenfeld, A. (1983). A theory of textural segmentation. In *Human and Machine Vision*, J. Beck, B. Hope and A. Rosenfeld, eds. (Academic Press, New York), pp. 1-38.
- Bullier, J., Hupé, J. M., James, A., and Girard, P. (1996). Functional interactions between areas V1 and V2 in the monkey. *Journal of Physiology (Paris)* **90**, 217-220.
- Calloway, E. M. and Katz, L. C. (1990). Emergence and refinement of clustered horizontal connections in cat striate cortex. *Journal of Neuroscience* **10**, 1134-1153.
- Cannon, M. W. and Fullenkamp, S. C. (1993). Spatial interactions in apparent contrast: Individual differences in enhancement and suppression effects. *Vision Research* **33**, 1685-1695.
- Carpenter, G. and Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing* **37**, 54-115.
- Carpenter, G. and Grossberg, S., eds. (1991). *Pattern Recognition by Self-Organizing Neural Networks* (M.I.T. Press, Cambridge).
- Carpenter, G. and Grossberg, S. (1993). Normal and amnesic learning, recognition, and memory by a neural model of cortico-hippocampal interactions. *Trends in Neurosciences* **16**, 131-137.
- Cauler, L. J. and Connors, B. W. (1994). Synaptic physiology of horizontal afferents to layer I in slices of rat SI cortex. *Journal of Neuroscience* **14**, 751-762.
- Chapman, B. Zahs, K. R., and Stryker, M. P. (1991). Relation of cortical cell orientation selectivity to alignment of receptive fields of the geniculocortical afferents that arborize within a single orientation column in ferret visual cortex. *The Journal of Neuroscience* **11**, 1347-1358.
- Das, A. and Gilbert, C. D. (1995) Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature* **375**, 780-784.
- DeAngelis, G. C., Ohzawa, I. and Freeman, R. D. (1993). Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. 1. General characteristics and postnatal development. *Journal of Neurophysiology* **69**, 1091-1117.
- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A. C., and Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science* **269**, 981-985.
- Dubin, M. W. and Cleland, B. G. (1977). Organization of visual inputs to interneurons of lateral geniculate nucleus of the cat. *Journal of Neurophysiology* **40**, 410-427.

- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General* **113**, 501-517.
- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., and Reitbock, H. J. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics* **60**, 121-130.
- Felleman, D. J. and van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* **1**, 1-47.
- Ferster, D. and Lindström, S. (1983). An intracellular analysis of geniculo-cortical connectivity of area 17 of the cat. *Journal of Physiology* **342**, 181-215.
- Ferster, D. and Lindström, S. (1985). Synaptic excitation of neurones in area 17 of the cat by intracortical axon collaterals of cortico-geniculate cells. *Journal of Physiology* **367**, 233-252.
- Francis, G. and Grossberg, S. (1996a). Cortical dynamics of form and motion integration: Persistence, apparent motion, and illusory contours. *Vision Research* **35**, 149-173.
- Francis, G. and Grossberg, S. (1996b). Cortical dynamics of boundary segmentation and reset: Persistence, afterimages, and residual traces. *Perception* **25**, 543-567.
- Francis, G., Grossberg, S., and Mingolla, E. (1994). Cortical dynamics of feature binding and reset: Control of visual persistence. *Vision Research* **34**, 1089-1104.
- Galuske, R. A. W. and Singer, W. (1996). The origin and topography of long-range intrinsic projections in cat visual cortex: A developmental study. *Cerebral Cortex* **6**, 417-430.
- Ghose, G. M., Freeman, R. D., and Ohzawa, I. (1994). Local intracortical connections in the cat's visual cortex: Postnatal development and plasticity. *Journal of Neurophysiology* **72**, 1290-1303.
- Gilbert, C. D. and Wiesel, T. N. (1979). Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature* **280**, 120-125.
- Gilbert, C. D. and Wiesel, T. N. (1992). Receptive field dynamics in adult primary visual cortex. *Nature* **356**, 150-152.
- Grieve, K. L. and Sillito, A. M. (1991). The length summation properties of layer VI cells in the visual cortex and hypercomplex cell end zone inhibition. *Experimental Brain Research* **84**, 319-325.
- Grieve, K. L. and Sillito, A. M. (1991). A re-appraisal of the role of layer VI of the visual cortex in the generation of cortical end inhibition. *Experimental Brain Research* **87**, 521-529.
- Grieve, K. L. and Sillito, A. M. (1995). Non-length-tuned cells in layers II/III and IV of the visual cortex: The effect of blockade of layer VI on responses to stimuli of different lengths. *Experimental Brain Research* **104**, 12-20.
- Gove, A., Grossberg, S., and Mingolla, E. (1995). Brightness, perception, illusory contours, and corticogeniculate feedback. *Visual Neuroscience* **12**, 1027-1052.
- Gray, C. M. and Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences* **86**, 1698-1702.
- Grosovsky, D. H., Shapley, R. M., and Hawken, M. J. (1993). Macaque V1 neurons can signal "illusory" contours. *Nature* **365**, 550-552.
- Grossberg, S. (1969). On the production and release of chemical transmitters and related topics in cellular control. *Journal of Theoretical Biology* **22**, 325-364.
- Grossberg, S. (1976a). Adaptive pattern classification and universal recoding, I: Parallel development and coding of neural feature detectors. *Biological Cybernetics* **21**, 117-158.
- Grossberg, S. (1976b). Adaptive pattern classification and universal recoding, II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics* **23**, 187-202.

- Grossberg, S. (1973). Contour enhancement, short-term memory, and constancies in reverberating neural networks. *Studies in Applied Mathematics* **52**, 217-257.
- Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review* **87**, 1-51.
- Grossberg, S. (1982). *Studies of mind and brain*. Amsterdam: Kluwer.
- Grossberg, S. (1994). 3-D vision and figure-ground separation by visual cortex. *Perception and Psychophysics* **55**, 48-120.
- Grossberg, S. (1995). The attentive brain. *American Scientist* **83**, 483-449.
- Grossberg, S., Boardman, I., and Cohen, M. (1997). Neural dynamics of variable-rate speech categorization. *Journal of Experimental Psychology* **2**, 481-503.
- Grossberg, S. and Merrill, J. W. L. (1996). The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. *Journal of Cognitive Neuroscience* **8**, 257-277.
- Grossberg, S. and Mingola, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception and Psychophysics* **38**, 141-171.
- Grossberg, S., Mingolla, E., and Ross, W. D. (1997). Visual brain and visual perception: How does the cortex do perceptual grouping? *Trends in Neurosciences* **20**, 106-111.
- Grossberg, S. and Somers, D. (1991). Synchronized oscillations during cooperative feature linking in a cortical model of visual perception. *Neural Networks* **4**, 453-466.
- Grossberg, S. and Stone, G. O. (1986). Neural dynamics of word recognition and recall: Attentional priming, learning, and resonance. *Psychological Review* **93**, 46-74.
- Grossberg, S. and Williamson, J. R. (1997). Linking cortical development to visual perception. *Society for Neuroscience Abstracts* **227.9**, 568.
- Hirsch, J. A. and Gilbert, C. D. (1991). Synaptic physiology of horizontal connections in the cat visual cortex. *The Journal of Neuroscience* **11**, 1800-1809.
- Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology* **160**, 106-154.
- Hubel, D. H., Wiesel, T. N., and LeVay, S. (1977). Plasticity of ocular dominance columns in monkey striate cortex. *Philosophical Transactions of the Royal Society of London (B)* **278**, 377-409.
- Hubel, D. H. and Wiesel, T. N. (1977). Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society of London* **198**, 1-59.
- Hupé, J. M., James, A. C., Girard, P., and Bullier, J. (1997). Feedback connections from V2 modulate intrinsic connectivity within. *Society for Neuroscience Abstracts* **406.15**, 1031.
- Ito, M., Westheimer, G., and Gilbert, D. C. (1997). Attention modulates the influence of context on spatial integration in V1 of alert monkeys. *Society for Neuroscience Abstracts* **603.2**, 1031.
- Johnson, R. R. and Burkhalter, A. (1997). A circuit for amplification of excitatory feedback input from rat extrastriate cortex to primary visual cortex. *Society for Neuroscience Abstracts* **651.7**, 1669.
- Julesz, B. (1971). *Foundations of Cyclopean Perception* (University of Chicago Press, Chicago).
- Kandel, E. R. and O'Dell, T. J. (1992). Are adult learning mechanisms also used for development? *Science* **258**, 243-245.
- Karni, A. and Sagi, D. (1991). Where practice makes perfect in textural discrimination: Evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences* **88**, 4966-4970.
- Kisvárdy, Z. F., Tóth, E., Rausch, M., and Eysel, U. T. (1995). Comparison of lateral excitatory and inhibitory connections in cortical orientation maps of the cat. *Society for Neuroscience Abstracts* **360.11**, 907.

- Knierim, J. J. and van Essen, D. C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *Journal of Neurophysiology* **67**, 961-980.
- Lamme, V. A. F., Zipser, K., and Spekreijse, H. (1997). Figure-ground signals in V1 depend on consciousness and feedback from extra-striate areas. *Society for Neuroscience Abstracts* **603.1**, 1543.
- Lowel, S. and Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science* **255**, 209-212.
- Luhmann, H. J., Martinez Millan, L., and Singer, W. (1986). Development of horizontal intrinsic connections in cat striate cortex. *Experimental Brain Research* **63**, 443-448.
- Macchi, G. and Rinvik, E. (1976). Thalamo-telencephalic circuits: A neuroanatomical survey. In A. Remond Ed. *Handbook of Electroencephalography and Clinical Neurophysiology Vol. 2*, Pt. A, (Elsevier, Amsterdam).
- Mayford, M., Barzilai, A., Keller, F., Schacher, S., and Kandel, E. R. (1992). Modulation of an NCAM-related adhesion molecule with long-term synaptic plasticity in aplysia. *Science* **256**, 638-644.
- McAdams, C. J. and Maunsell, J. H. R. (1997). Spatial attention and feature-directed attention can both modulate neuronal responses in macaque area V4. *Society for Neuroscience Abstracts* **802.5**, 2062.
- McGuire, B. A., Gilbert, C. D., Riolin, P. A., and Wiesel, T. N. (1991) Target of horizontal connections in macaque primary visual cortex. *The Journal of Comparative Neurology* **305**, 370-392.
- Merzenich, M. M., Recanzone, E. G., Jenkins, W. M., Allard, T. T., and Nudo, R. J. (1988). in *Neurobiology of Neocortex*, P. Rakic and W. Singer, eds. (Wiley, New York), pp. 41-67.
- Motter, B. C. (1994a). Neural correlates of attentive selection for color or luminance in extrastriate area V4. *Journal of Neuroscience* **14**, 2178-2189.
- Motter, B. C. (1994b). Neural correlates of feature selective memory and pop-out in extrastriate area V4. *Journal of Neuroscience* **14**, 2190-2199.
- Mountcastle, V. B. (1957). Modality and topographic properties of single neurons of cats somatic sensory cortex. *Journal of Neurophysiology* **20**, 408-434.
- Murphy, P. C. and Sillito, A. M. (1987). Cortico fugal feedback influences the generation of length tuning in the visual pathway. *Nature* **329**, 727-729.
- Murphy, P. C. and Sillito, A. M. (1996). Functional morphology of the feedback pathway from area 17 of the cat visual cortex to the lateral geniculate nucleus. *Journal of Neuroscience* **16**, 1180-1192.
- Meyer, G., Lawson, R., and Cohen, W. (1975). The effects of orientation-specific adaptation on the duration of short-term visual storage. *Vision Research* **15**, 569-572.
- O'Craven, K. M., Rosen, B. R., Kwong, K. K., Treisman, A., and Savoy, R. L. (1997). Voluntary attention modulates fMRI activity in human MT-MST. *Neuron* **18**, 591-598.
- Pandya, D. N. and Yeterian, E. H. (1985). Architecture and connections of cortical association areas. In A. Peters and E. G. Jones, eds. *Cerebral Cortex* **10** (Plenum Press, New York).
- Peterhans, E. and von der Heydt (1989). Mechanisms of contour perception in monkey visual cortex. II. Contours bridging gaps. *The Journal of Neuroscience* **9**, 1749-1763.
- Poggio, T., Fahle, M., and Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science* **256**, 1018-1021.
- Polat, U. and Sagi, D. (1994). The architecture of perceptual spatial interactions. *Vision Research* **34**, 73-78.

- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology* **32**, 3-25.
- Press, W. A. and van Essen, D. C. (1997). Attentional modulation of neuronal responses in macaque area V1. *Society for Neuroscience Abstracts* **405.3**, 1026.
- Ramachandran, V. S. and Nelson, J. I. (1976). Global grouping overrides point-to-point disparities. *Perception* **5**, 125-128.
- Redies, C. Crook, J. M., and Creutzfeldt, O. D. (1986). Neural responses to borders with and without luminance gradients in cat visual cortex and dLGN. *Experimental Brain Research* **61**, 469-481.
- Reid, R. C. and Alonso, J-M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* **378**, 281-284.
- Reynolds, J., Nicholas, J., Chelazzi, L., and Desimone, R. (1995). Spatial attention protects macaque V2 and V4 cells from the influence of non-attended stimuli. *Society for Neuroscience Abstracts* **21.3**, 1759.
- Rockland, K. S. (1994). The organization of feedback connections from area V1(18) to V1(17). In A. Peters and K. S. Rockland, eds. *Cerebral Cortex* **4** (Plenum Press, New York).
- Schmidt, K. E., Schlote, W. Bratzke, H., Rauen, T., Singer, W., and Galuske, R. A. W. (1997). Patterns of long range intrinsic connectivity in auditory and language areas of the human temporal cortex. *Society for Neuroscience Abstracts* **415.13**, 1058.
- Sillito, A. M., Jones, H. E., Gerstein, G. L., and West, D. C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature* **369**, 479-482.
- Singer, W. (1982). The role of attention in developmental plasticity. *Human Neurobiology* **1**, 41-43.
- Somers, D. C., Nelson, S. B. and Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience* **15**, 5448-5465.
- Stemmler, M., Usher, M., and Niebur, E. (1995). Lateral interactions in primary visual cortex: A model bridging physiology and psychophysics. *Science* **269**, 1877-1880.
- Stryker, M. P. and Harris, W. (1986). Binocular impulse blockade prevents the formation of ocular dominance columns in cat visual cortex. *Journal of Neuroscience* **6**, 2117-2133.
- Treue, S. and Maunsell, J. H. R. (1997). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* **382**, 539-541.
- Tsumoto, T., Creutzfeldt, O. D., and Legendy, C. R. (1978). Functional organization of the corticofugal system from visual cortex to lateral geniculate nucleus in the cat. *Experimental Brain Research* **25**, 291-306.
- van Essen, D. C. and Maunsell, J. H. R. (1983). Hierarchical organization and functional streams in the visual cortex. *Trends in Neurosciences* **6**, 370-375.
- von der Heydt, R. Peterhans, E., and Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science* **224**, 1260-1262.
- Weber, A. J., Kalil, R. E., and Behan, M. (1989). Synaptic connections between corticogeniculate axons and interneurons in the dorsal lateral geniculate nucleus of the cat. *Journal of Comparative Neurology* **289**, 156-164.
- Wittmer, L. L., Dalva, M. B., and Katz, L. C. (1997). Reciprocal interactions between layer 4 and layer 6 cells in ferret visual cortex. *Society for Neuroscience Abstracts* **651.5**, 1668.
- Zeki, S. and Shipp, S. (1988). The functional logic of cortical connections. *Nature* **335**, 311-317.
- Zohary, E., Cerebrini, S., Britten, K. H., and Newsome, W. T. (1994). Neuronal plasticity that underlies improvement in perceptual performance. *Science* **263**, 1289-1292.