

Directed Visual Attention and the Dynamic Control of Information Flow

Charles H. Anderson, David C. Van Essen^{a,*}
Bruno A. Olshausen^b

^a*Washington University School of Medicine*

^b*Center for Neuroscience, UC Davis
and Redwood Neuroscience Institute*

Abstract

Visual attention serves as an information bottleneck, allowing for efficient analysis of a region of interest that shifts in location and spatial scale from moment to moment. We discuss a modeling framework for visual attention in which information flow through the visual hierarchy is regulated by dynamic control of connection strengths. A key aspect of this model involves the establishment of object-centered reference frames for visual working memory as well as object recognition. Psychophysical evidence suggests that the region of interest is about 30 resolution elements across. This model is neurobiologically plausible, supported by several lines of anatomical and physiological data, and suitable for embedding into a larger computational framework for modeling of neural representations and transformations.

Key words: visual attention, dynamic routing, invariance

* Dept. of Anatomy & Neurobiology
Campus Box 8108
660 S. Euclid
St. Louis, MO 63110-1093

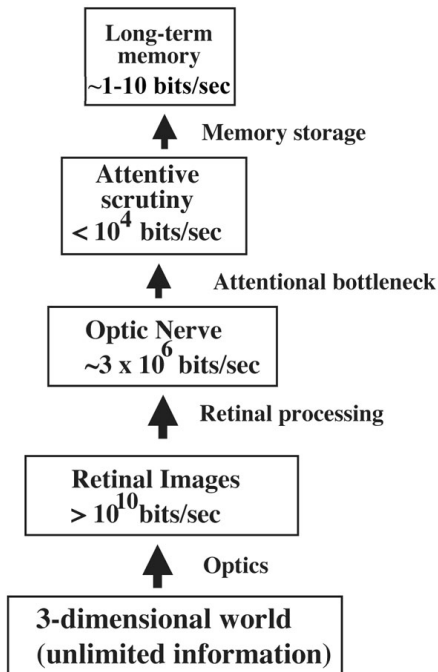
Phone: 314/362-7043
FAX: 314/747-3436
vanessen@v1.wustl.edu

1 Introduction

Neural encoding of retinal images generates a barrage of information transmitted along the optic nerves that far exceeds what can actually be perceived from moment to moment. Visual attention is a process that addresses this problem by directing a tiny fraction of the information arriving at primary visual cortex to high-level centers involved in visual working memory and pattern recognition. In the information pyramid shown in Figure 1, visual attention constitutes a major bottleneck. Much less than one percent of the information carried through the optic nerves reaches attentive scrutiny (see figure legend for details). Because the attended information comes from a spatially restricted region of interest (ROI), visual attention must mediate translation, scale and other invariances, so that under a variety of different poses, a given object or spatial pattern can be recognized and analyzed for how it might be manipulated.

A computationally sound and neurobiologically plausible model of visual attention should address five broad issues, reflecting the reality that attention is an extremely sophisticated process engaging much if not most of the entire visual system. (i) What are the high-level brain structures that mediate visual pattern recognition and working memory? (ii) What is the anatomical circuitry that provides ascending visual inputs to these high-level centers? (iii) What are the control structures that determine where attention should be directed from moment to moment? (iv) How do these control signals dynamically modulate the connections leading up to high-level centers and thereby regulate information flow? (v) How do non-attentive visual processes that operate on the whole image support attentive visual processing?

In previous publications, we have presented a conceptual and computational framework for modeling visual attention as it relates to object recognition (Anderson & Van Essen, 1987; Olshausen et al., 1993, 1995; Van Essen et al., 1994). Central to this framework are the notions that (i) information flow into inferotemporal cortex is controlled dynamically across multiple hierarchical levels; (ii) spatial relationships within the attended ROI are preserved explicitly; and (iii) attentional control operates by dynamic modulation of connectivity weights. Here, we reassess and refine these core concepts and extend them to incorporate visual working memory as well as object recognition. We also describe initial steps towards integrating attention into a broader computational framework for modeling neural representations and transformations (Eliasmith & Anderson, 2003).



AN INFORMATION PYRAMID

Fig. 1. An information pyramid for the visual system. The estimates of information available in retinal images, encoded in the optic nerves, and passed through the window of attention have been discussed by Van Essen et al. (1991) and Van Essen & Anderson (1995) and are based on information rates of 3 bits/s for each neuron (Eliasmith & Anderson, 2003). Human memory storage rates are based on estimates by Landauer (1986).

2 Dynamic routing

Dynamic control of information flow is fundamental to how both brains and modern digital computers deal with finite resources, whether it be for number-crunching in a computer or for sensory processing, motor control, or cognitive processing in the brain. In computers, routing circuits constitute major components of the hardware, and a large fraction of the software is devoted to the issue of where to get information for the task at hand and where to put it after it has been processed. In the brain, dynamic reconfiguration of neural resources must be done through modulation of connectivity. We refer to this as *dynamic routing* due to the close analogy with computer routing circuits. A major difference between brains and digital computers is that memory and processing are tightly bound to one another in the brain, whereas in computers they are distinct. In essence, the brain contains numerous interlinked subsystems, each carrying out a set of subroutines for specialized computations that would be handled generically by a computer CPU.

In the visual cortex, dynamic routing is most efficiently accomplished in a hierarchical fashion, starting with local control of information flow in lower visual areas and extending to more global control at higher levels. Control signals can also operate both top down and bottom up. Attention is normally associated with global top-down processes that are accessible through conscious awareness, though bottom-up control signals can be regarded as involuntary attention (Yantis & Jonides, 1984). Local control interactions are not generally considered attention, although they may involve similar operations at the neural level. Overall, ‘attention’ is inherently an imprecise term that captures only a segment of the broader and more fundamental issue of the control of information flow.

We believe that dynamic routing is essential for achieving invariant representations, both for object recognition and for working memory. An important part of achieving invariant representations is to preserve information about the variations themselves. That is, *in order to interpret what has not changed, it is often important also to know what has changed*. Dynamic routing accomplishes this by remapping topographic representations from one level to the next. The control signals mediating the remapping contain information about the variations, while the remapped information itself constitutes the invariant part (see also Arathorn chapter). This is in sharp contrast to the many alternative models of object recognition which do not attempt to model the variations but instead attempt to compute collections of loosely assembled “invariant features” (Selfridge, 1959; Fukushima, 1980; LeCun et al., 1990; Riesenhuber & Poggio, 1999). While such Pandemonium-style models may be capable of simple object discriminations, we contend that they are insufficient to account for more general aspects of perception and visuo-motor function in the real world. Examples include the perception of generic surfaces (a crumpled cloth); the ability not only to read hand-written text but also to recognize the idiosyncratic style of the writer; and the ability to geometrically manipulate objects in working memory and plan appropriate grasping maneuvers.

In this chapter we focus on routing circuits for translation and scale invariance, as these are fundamental requirements of any flexible vision system. Scale invariance imposes particularly stringent constraints on system design for human vision, given that we can recognize a face when it is only a foot away, where the image covers most of the visual field and occupies most of primary visual cortex (~ 40 cm² in both hemispheres), or when it is at a distance of 100 feet, where the image subtends about 0.4 degrees and occupies only a few percent of primary visual cortex (~ 1 cm²).





Fig. 2. Dynamic routing circuit. *top*, Connections from one level to the next are dynamically modulated by a set of control neurons. *bottom*, Shown are two different states of effective connectivity, corresponding to different sizes and positions of the ROI.

3 Dynamic routing circuit architecture

The basic architecture we propose for dynamic routing in the visual cortex is schematized in Figure 2. The input layer at the bottom of the array contains the inputs from the retina. The diagram illustrates a spatially uniform sampling, though in fact the distribution of retinal ganglion cell inputs declines steeply with retinal eccentricity (see Olshausen et al. (1995), for a routing circuit that utilizes a retinal sampling lattice). The top layer represents many fewer spatial locations, reflecting the limited resolution within the ROI (see below). This upper layer may represent cortical areas involved in visual working memory and/or areas involved with object recognition. Note that no attempt is made to specify the features represented by these nodes, as there is very little neurophysiological data at present suggesting what these features should be. For now each node may be thought of as a feature vector.

The exponential increase in the spread of the receptive fields at each level and the subsampling at higher levels are essential elements of this circuit. This design minimizes the number of levels and reduces the complexity of the control, while respecting the fixed number of fan-in/fan-out connections between neurons. Many aspects of this design can be seen in the anatomy of the connections between the visual areas in the V1-IT pathway. For example the size of V1 and V2 are approximately the same, while V4 is significantly smaller. Also, the size of the axonal spread increases as one goes up to higher levels, and nominal receptive field sizes increase exponentially as well.

Topographic representations are remapped from one level to the next by a set of control signals that modulate the connections between each level through multiplicative interactions. Formally, the visual data signals at node μ and level l , X_μ^l , are computed by combining the visual data signals at the nodes in the layer below, X_μ^{l-1} , and the control signals at level l , C_λ^l , via

$$X_\mu^l = \sum_{\lambda, \nu} C_\lambda^l \Gamma_{\lambda\mu\nu}^l X_\nu^{l-1} \quad (1)$$

In order to see how certain geometric remappings could be accomplished by this circuit, it is helpful to rewrite Eq. (1) as follows

$$X_\mu^l = \sum_\nu W_{\mu\nu}^l X_\nu^{l-1} \quad (2)$$

$$W_{\mu\nu}^l = \sum_\lambda C_\lambda^l \Gamma_{\lambda\mu\nu}^l \quad (3)$$

A spatial pattern in level $l - 1$ may be translated by an amount d^l and scaled by an amount α^l in level l by setting the coupling matrix as follows

$$W_{\mu\nu}^l = \exp \left[-\frac{1}{2} \left(\frac{\mu - \alpha^l(\nu - d^l)}{\sigma} \right)^2 \right] \quad (4)$$

The problem of achieving a certain remapping d^l, α^l , at level l thus amounts to carving out the appropriate pattern in connection space (μ, ν) . In general this may be accomplished by viewing the three-way coupling constants $\Gamma_{\lambda\mu\nu}^l$ as basis functions $\Psi_\lambda^l(\mu, \nu)$. Setting the control signals appropriately then amounts to a function approximation problem:

$$\exp \left[-\frac{1}{2} \left(\frac{\mu - \alpha^l(\nu - d^l)}{\sigma} \right)^2 \right] = \sum_\lambda C_\lambda^l(\alpha^l, d^l) \Psi_\lambda^l(\mu, \nu) \quad (5)$$

and the problem of learning amounts to finding a set of basis functions that produce good approximations for the ensemble of desired remappings.

Our previous proposal for implementing the multiplicative interactions in Eq. (1) was centered on pairwise, ‘and’-like interactions within dendritic trees. We now have a more general perspective based on the population coding framework described by Eliasmith & Anderson (2003). In this framework, a distinction is made between explicit neural activities and the implicit variables they represent. In the case of our routing circuit, the implicit variables are the image data signals, X_μ^l , which are represented by neural activities a_i^l in terms

of a linear basis function expansion

$$\mathbf{X}^l = \sum_i \phi_i^{\mathbf{X}} a_i^l(\mathbf{X}^{l-1}, \mathbf{C}^l) \approx \sum_\lambda \mathbf{C}_\lambda^l \Gamma_\lambda^l \cdot \mathbf{X}^{l-1} \quad (6)$$

where the boldface variables represent vectors, and the basis vectors $\phi_i^{\mathbf{X}}$ have the same dimensionality as \mathbf{X}^l . The neuronal activities are computed via

$$a_i^l(\mathbf{X}^{l-1}, \mathbf{C}^l) = \mathcal{G} \left[\tilde{\phi}_i^{\mathbf{X}} \cdot \mathbf{X}^{l-1} + \tilde{\phi}_i^{\mathbf{C}} \cdot \mathbf{C}^l \right] \quad (7)$$

where $\mathcal{G}[\]$ denotes the integrate-and-fire response nonlinearity of a neuron. The input consists of the sum of a linear projection of the data signal space at the level below and a projection of the control neuron activities. Given a population of neurons a_i^l that represent the joint space $(\mathbf{X}^{l-1}, \mathbf{C}^l)$ with sufficient numbers and diversity, one can obtain the transformation described in Eq. (6). Methods for setting the parameters $\tilde{\phi}_i^{\mathbf{X}}, \phi_i^{\mathbf{X}}, \tilde{\phi}_k^{\mathbf{C}}$ and $\phi_k^{\mathbf{C}}$ are described in Eliasmith & Anderson (2003).

The neuronal activities corresponding to the next level are similarly computed via

$$a_i^{l+1}(\mathbf{X}^l, \mathbf{C}^{l+1}) = \mathcal{G} \left[\tilde{\phi}_i^{\mathbf{X}} \cdot \mathbf{X}^l + \tilde{\phi}_i^{\mathbf{C}} \cdot \mathbf{C}^{l+1} \right] \quad (8)$$

Expanding \mathbf{X}^l as in Eq. (6) yields

$$a_j^{l+1}(\mathbf{X}^l, \mathbf{C}^{l+1}) = \mathcal{G} \left[\sum_i \tilde{\phi}_j^{\mathbf{X}} \phi_i^{\mathbf{X}} a_i^l(\mathbf{X}^{l-1}, \mathbf{C}^l) + \tilde{\phi}_j^{\mathbf{C}} \cdot \mathbf{C}^{l+1} \right] \quad (9)$$

The complete computation from level $l - 1$ to level $l + 1$ is illustrated in Figure 3.

Equation (9) together with Figure 3 illustrates that the input to the neurons in one layer consists of a weighted sum over units in the previous layer. Alternatively, the $a_i^l(\mathbf{X}^{l-1}, \mathbf{C}^l)$ in Eq. (9) may correspond to dendritic subunits of the level $l + 1$ neurons. In the latter case they act similarly to our previous suggestion of selective pruning in the dendrites, except that now only local summation within a nonlinear subunit is required rather than the precise pairwise multiplicative interactions in Eq. (1). Such nonlinear interactions could easily be realized by the biophysical properties of dendritic trees (Mel, 1997).




Fig. 3. Neuronal Circuit: The signals \mathbf{X}^{l-1} and \mathbf{C}^l are encoded into neural activities a_i^l using an additive projection rule. The multiplicative terms $\mathbf{C}_{\lambda}^l \Gamma_{\lambda}^l \mathbf{X}^{l-1}$ are then linearly decoded from this population of neurons.

4 Autonomous control

The control neurons are driven in one of two modes: 1) a bottom-up mode in which they receive input from a saliency map, or 2) a top-down mode in which they are driven by multiplicatively combining the output of memory with the input. In both cases, the dynamics of the control neurons are governed by an energy function that measures the correlation between the input nodes X_{μ}^{l-1} and a template V_{μ}^l as coupled through the control neurons, in addition to a constraint on the control neurons that encourages control states corresponding to affine transformations. For a single stage of the routing circuit this amounts to

$$E^l = - \sum_{\lambda, \mu, \nu} V_{\mu}^l C_{\lambda}^l \Gamma_{\lambda \mu \nu}^l X_{\nu}^{l-1} - \sum_{\lambda, \nu} C_{\lambda}^l T_{\lambda \nu}^l C_{\nu}^l \quad (10)$$

In the bottom-up case, V_{ν}^l is fixed to a Gaussian blob, whereas in the top-down case the V_{ν}^l correspond to the outputs of an associative memory (e.g., a Hopfield network). The control neurons are then computed by performing gradient descent on Eq. (10):

$$\begin{aligned} C_{\lambda}^l &= \sigma(u_{\lambda}^l) \\ \dot{u}_{\lambda}^l + \tau^{-1} u_{\lambda}^l &= \sum_{\mu, \nu} V_{\mu}^l \Gamma_{\lambda \mu \nu}^l X_{\mu}^{l-1} X_{\nu}^{l-1} + \beta \sum_{\nu} T_{\lambda \nu}^l C_{\nu}^l \end{aligned} \quad (11)$$

where σ is a sigmoidal nonlinearity. The control neurons are also subject to delayed self-inhibition so that the circuit does not perseverate on any one ROI. The full control sequence proceeds as follows:

- (1) set V_{μ}^l to a Gaussian blob (focus ROI on something interesting),

- (2) switch V_μ^l to the output associative memory, and
- (3) inhibit the current control state and return to (1).

When an object is being recognized then the V_μ^l evolve according to the recall dynamics of memory. Alternatively if an object is being searched for, then the V_μ^l may be set to correspond to the object of interest, and attention will then be directed to those locations of the image containing the object. This mode would most closely correspond to “object-based attention.”

5 Neurobiological substrates and mechanisms

We have shown previously (Olshausen et al., 1993, 1995) that the routing circuit of Figure 2 may be scaled up in a manner consistent with the anatomy of the ventral stream in visual cortex ($V1 \rightarrow V2 \rightarrow V4 \rightarrow IT$). Such a circuit would comprise approximately 300,000 sample nodes (in 2D) at the input in V1, narrowing to approximately 1000 sample nodes in IT. The pulvinar is hypothesized to be a major source of the control signals modulating connectivity along the ventral stream, although control neurons could reside in the cortex as well.

There are two potential candidates for the top ROI level, which are not mutually exclusive. One is dedicated to the task of object recognition, which most likely resides in inferotemporal cortex. The other provides the substrate for visual working memory (VWM), which probably resides in parietal and/or frontal cortex. In VWM the image data is utilized for planning how to manipulate the objects in the external environment, where it is clear local spatial relationships must be made available. Several lines of psychophysical experiments suggest that VWM is limited to representing something like a region of space with a diameter of ~ 30 distinct spatial locations at any instant in time (Van Essen et al., 1991). A circuit like the one outlined above would allow the visual system to dynamically rescale and reposition the part of the visual field that gets mapped into VWM. Our current conception of VWM is that it is a distinct and separate entity from the cortical areas involved in object recognition, although the circuits for mediating both are closely related, or may in fact be the same. New objects are scrutinized using VWM, but recognition is transferred to the object recognition area after the system becomes familiar with the object. In doing so, aspects of the 30×30 size of VWM and details of local spatial relationships can be maintained, but need not be. Thus, rapid estimates of what objects are in the visual field may take place without these details, but when the task becomes more complex the details of local spatial relationships become very important.

5.1 *Physiological evidence*

The model is consistent with a number of attentional effects seen in visual cortex, most notably receptive field shifts and gain modulation. Neurons at intermediate stages of the circuit would be expected to show pronounced receptive field shifts as attention is moved to different locations in the visual field, as shown in Figure 1. Such shifts are evident in a number of neurons in V4 as reported by Connor et al. (1997). However, gain changes would also be expected, since according to the population coding framework above the neuronal activities would be subject to modulation by control neurons (Eq. 7). Such gain changes were also reported by Connor et al. (1997).

The model further stipulates that the beginnings of an object-centered reference frame should be established at intermediate or high levels of the cortical hierarchy. Several neurophysiological studies provide strong evidence for object-centered processing in area V4. Connor et al. (1997) demonstrated that responses to a given stimulus within a cell’s receptive field can be dramatically modulated according to whether attention is directed to one side or another of the receptive field. Pasupathy & Connor (2001) demonstrated that many V4 neurons encode shape characteristics within restricted subregions of a larger object. In inferotemporal cortex, there is evidence for size-invariant tuning in many neurons. It has yet to be determined whether there is explicit coding of location within the attended region in inferotemporal cortex or in frontal and parietal areas engaged in visual working memory.

Most other neurophysiological studies of attention in V4 and other visual areas have focused on different aspects of attention (e.g., the degree of enhancement produced by attention to a given stimulus). The attentional modulation reported in these studies is generally less pronounced than that found by Connor et al. (1997). Thus, while much remains to be determined about the physiological basis of attentional processing, there is solid support for the general hypothesis of object-centered processing established across multiple stages of the cortical hierarchy.

5.2 *Experimental Predictions*

One of the main predictions of the model is that there exists a separate and distinct population of neurons—control neurons—whose job it is to route information flow in the cortex. The activity of such neurons would reflect solely the attentional state rather than the contents of the ROI per se. However, at lower levels of the cortical hierarchy the control neurons will not necessarily be tied to overt attention (as discussed previously), and so they may be difficult

to correlate with behavior.

Another major prediction of the model is that invariance is mediated by dynamic routing, rather than through a feedforward filter bank of invariant features as in a Pandemonium-style model. As a consequence, one should expect to find congruency effects in object recognition—i.e., recognizing an object at one position and size should make it easier to recognize a subsequent (different) object at the same position and size because the ROI would not need to be changed. While there is already some evidence for such effects in recognition, additional experiments are needed to disambiguate what exactly is being primed in the system (features at the same location vs. the state of control). A related prediction is that the ROI should be limited to a fixed resolution ($\sim 30 \times 30$ sample nodes). Previously we have documented several lines of evidence that suggest this idea (Van Essen et al., 1991), but the current state of affairs is still inconclusive and deserves further investigation.

6 Discussion

An important attribute of any model of brain function is that it not only be able to fit the available anatomical and physiological data, but also that it is capable of explaining function. The dynamic routing model of attention was designed to solve an important problem of vision - invariant object representation. In that regard it is distinct from many other models that focus mainly on emulating neural response properties or aspects of psychophysical performance measured in laboratory settings. It also differs from many other models of object recognition in that it represents both the *variations* and the *invariances* of objects, rather than simply trying to capture the invariant part as in Pandemonium-style models.

Since our model was originally proposed, several other investigators have developed alternative models for achieving invariant representations that also draw upon the idea of dynamic routing. Salinas & Abbott (1997) have proposed a model of area V4 that utilizes “gain fields” to transform reference frames and is similar to our development above using the population coding framework, although their model utilizes direct multiplication on each neuron rather than sigmoidal nonlinearities. Amit & Mascaró (2003) have proposed a hierarchical model of recognition that uses a replica module to represent many different shifted or rescaled versions of the input and then switches between them. This may be seen as a special case of our routing circuit in which the control basis functions correspond to entire shifts or rescalings of a large region of the image. Arathorn (2002; see also chapter XX in this volume) has proposed a model very similar to ours that utilizes multiplicative gating on the inputs, and he has some very compelling demonstrations of its effectiveness in

real-world object recognition tasks.

Our proposal remains a zeroth-order model in the sense that there are still many components that require better specification in order to offer more detailed experimental predictions. Chief among these is the choice of feature representation—i.e., what is actually being represented by the sample nodes in each area? A promising approach to this question is to use unsupervised learning procedures to discover what forms of structure are best suited for representing natural images. Exploring such procedures for learning the features, in addition to the control neuron weights, remains the subject of current investigation.

References

- Amit, Y. and Mascaró, M. (2003). An integrated network for invariant visual detection and recognition, *Vision Research*, 43, 2073-2088.
- Anderson, C.H. and Van Essen, D.C. (1987). Shifter circuits: a computational strategy for dynamic aspects of visual attention. *Proc. Natl. Acad. Sci.* 84, 6297-6301.
- Arathorn, David W. (2002). *Map-Seeking Circuits in Visual Cognition*, Stanford University Press.
- Connor, C.E., Preddie, D.G., Gallant, J.L. and Van Essen, D.C. (1997). Spatial attention effects in macaque area V4. *J. Neurosci.* 17, 3201-3214.
- Eliasmith, C. and Anderson C.H. (2003). *Neural Engineering*. MIT Press.
- Fukushima, K. (1980). Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36, 193-202.
- Landauer, T.K. (1986) How much do people remember? Some estimates of the quantity of learned information in long-term memory. *Cognitive Science* 10, 477-493, 1986.
- LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., Jackel, L.D. (1990). Backpropagation applied to handwritten Zip code recognition. *Neural Computation*, 1, 541-551.
- Mel, B.W. (1994) Information processing in dendritic trees. *Neural Computation*, 6, 1031-1085.
- Olshausen, B.A., Anderson, C.H., and Van Essen, D.C. (1993). A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *J. Neurosci.* 13, 4700-4719.
- Olshausen, B.A., Anderson, C.H. and Van Essen, D.C. (1995). A multiscale dynamic routing circuit for forming size- and position-invariant object representations. *J. Computational Neuroscience* 2, 45-62.
- Pasupathy, A. and Connor, C.E. (2001). Shape Representation in area V4:

- position-specific tuning for boundary conformation. *J. Neurophysiol.* 86, 2505-2519.
- Riesenhuber, M., and Poggio, T. (1999). Hierarchical Models of Object Recognition in Cortex. *Nature Neuroscience* 2, 1019-1025.
- Selfridge, O.G. (1959). Pandemonium: a paradigm for learning. In: *The Mechanisation of Thought Processes*, pp. 511-527. H.M.S.O., London, 1959.
- Salinas, E., and Abbott, L.F. (1997). Invariant visual responses from attentional gain fields. *Journal of Neurophysiology* 77, 3267-3272.
- Van Essen, D.C., Olshausen, B., Anderson, C.H. and Gallant, J.L. (1991). Pattern recognition, attention, and information bottlenecks in the primate visual system. In: *Proc. SPIE Conf. on Visual Information Processing: From Neurons to Chips, vol. 1473*, 17-28.
- Van Essen, D.C., Anderson, C.H., and Olshausen, B.A. (1994). Dynamic routing strategies in sensory, motor, and cognitive processing. In: *Large Scale Neuronal Theories of the Brain* (Eds. C. Koch and J. Davis), MIT Press, pp. 271-299.
- Van Essen, D.C. and Anderson, C.H. (1995) Information processing strategies and pathways in the primate visual system. In: *An Introduction to Neural and Electronic Networks*, 2nd ed., S. Zornetzer, J.L. Davis, C. Lau, T. McKenna, eds., Academic Press, FL, pp. 45-76.
- Yantis, S. and Jonides, J. (1984). Abrupt visual onsets and selective attention: Evidence from visual search. *Journal of Experimental Psychology: Human Perception & Performance*, 10, 601-621.