

# Triangles, Pyramids, Connections and Attentive Inhibition

**John K. Tsotsos**

Department of Computer Science  
University of Toronto  
Toronto, ONT M5S 3G4  
CANADA

tsotsos@cs.toronto.edu  
<http://www.cs.toronto.edu/~tsotsos>

Copyright (c) John K. Tsotsos 1999

PSYCHE, 5(20), July 1999  
<http://psyche.cs.monash.edu.au/v5/psyche-5-20-tsotsos.html>

KEYWORDS: vision, attention, inhibition, routing, visual pyramid, attentional control

COMMENTARY ON: LaBerge, D. (1997) Attention, Awareness, and the Triangular Circuit. *Consciousness and Cognition*, 6, 149-181. (See also LaBerge's precis for PSYCHE: <http://psyche.cs.monash.edu.au/v4/psyche-4-07-laberge.html>)

ABSTRACT: LaBerge's *Triangular Theory of Attention* contributes to several important topics in the study of visual attention. First, it expands on the discussion of whether attentive influences manifest themselves as neural suppression or enhancement; LaBerge seems to favour the enhancement viewpoint. Second, the paper proposes circuit loops (triangles by nature of three nodes in the loop) that may be responsible for the observed enhancement. Finally, a link between awareness and attention is explored and the claim that a representation of self must be considered is made. Here, it will not be possible to provide sensible discussion on all of these points; rather, the focus will be on the first issue, namely whether attention is manifested as enhancement or suppression. I claim that observations of enhancement or suppression depend very much on exactly how measurements are taken. Specific predictions are made: a given neuron exhibits enhanced or suppressed responses as a result of attentive influence depending on where that neuron is located in relation to the three-dimensional structure of attentive influences within the visual processing network. I show here a local, internal, attention control mechanism as an alternate model that provides an explanation for both enhancement and suppression of neural responses. In addition, it solves an important second problem for attention models, namely, information routing, within the same mechanism. The routing problem seems to be ignored by LaBerge.

# 1. Introduction: Enhancement or Suppression?

It has been a long-standing debate whether attention operates by re-inforcement, inhibition or a combination of inhibition and re-inforcement. LaBerge has a very nice diagram in his paper, Figure 1, that presents the options pictorially. In analyses previously presented (Tsotsos 1990, 1993, 1995a, 1995b) I have argued that inhibition of distractors is the correct choice among the options. Only in this way is the number of candidates to be considered in a search process reduced: the theory claims that the role of attention is to reduce the enormous number of possibilities that must be considered for solution to perception.

Inhibition of distractors also seems to be the only one of the choices which is sufficient to solve the routing problem, a problem not often discussed in the attention literature, but one which is critical to a successful realizable system. The routing problem, first brought to the fore by Anderson and Van Essen (1987), is this: as a signal traverses the visual processing network from the input to the output layers, it may be affected by signals from surrounding locations because of the convergence of information with the network. This effect, a context effect, may be detrimental to the veridical processing of the selected signal (see Connor et al., 1997). From a signal-to-noise perspective, the selected signal is optimally processed if it can traverse the visual processing network without interference. The obvious best strategy to achieve this would be to eliminate the effect of distractors.

Our model chooses to implement solutions to attentive selection and attentive routing by controlling whether connections, not neurons, are active or inhibited in the network. In this way, a given unit may participate in other computations while also having its effect on the computation of the attended signal minimized. This follows several previous works that also suggest connection inhibition (Hernandez-Peon, Scherrer, Jouvett, 1956; von der Malsburg, 1981; Tsotsos, 1991; Nyberg et al., 1996). This choice also has distinct implications for awareness (Tsotsos 1997). The next several sections will illustrate the model.

## 2. Attention and Pyramids

The visual processing architecture is assumed to be pyramidal in structure with units within this network receiving both feed-forward and feed-back connections (the model has this in common with the architecture developed in Van Essen et al. 1992). When a stimulus is first applied to the input layer of the pyramid, it activates in a feed-forward manner all of the units within the pyramid to which it is connected; the result is that an inverted sub-pyramid of units and connections is activated as shown in Figure 1A. It is assumed that response strength of units measures how strongly a neuron responds to the stimulus.

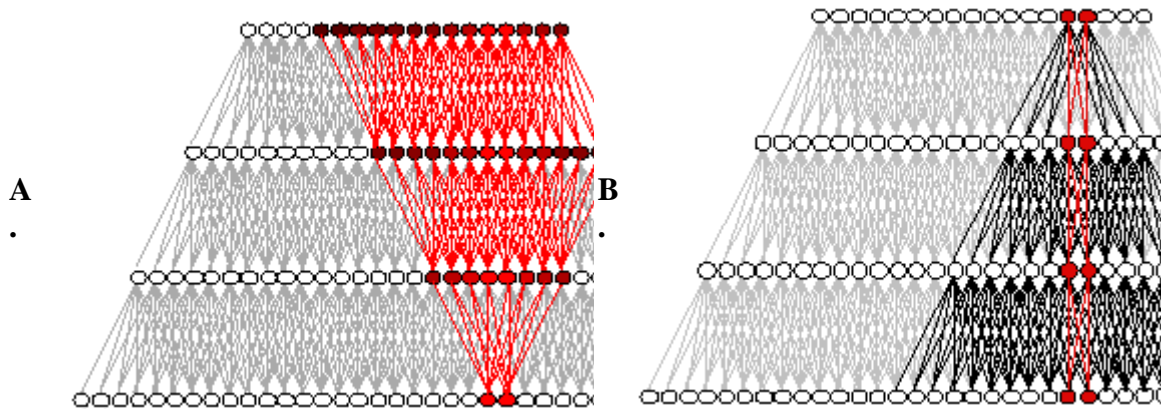
A single mechanism solves selection and minimization of interference of competing signals. Selection relies on a hierarchy of winner-take-all (WTA) processes. WTA is a parallel algorithm for finding the maximum value in a set of values. First, a WTA process operates across the entire visual field at the top layer: it computes the global winner, i.e., the units with largest response. The WTA can accept guidance for areas or stimulus qualities to favor if that guidance is available but operates independently otherwise. The search process then proceeds to the lower levels by activating a hierarchy of WTA processes. The global winner activates a WTA that operates only over its direct inputs. This localizes the largest response units within the top-level winning receptive field. Next, all of the connections of the visual pyramid that do not contribute to the winner are pruned. This strategy of finding the winners within successively smaller receptive fields, layer by layer in the pyramid and then pruning away irrelevant connections is applied recursively through the pyramid. The end result is that from the strongest response at the top of the network, the cause of that largest response is localized in the sensory field at the earliest levels. The paths remaining may be considered the pass zone while the pruned paths form the inhibitory zone of an attentional beam as in Figure 1B. The WTA does not violate biological connectivity or time constraints and the local inhibition required by the WTA is triggered by excitatory connections between areas.

The two examples which follow are intended to illustrate the structure and time course of the application of attentional selection in the model. The first example shown in Figure 1 shows a hypothetical visual processing pyramid; the figure caption provides further explanation. There are 4 layers, each unit connected to 7 units in the layer above it and 7 units in the layer below it. The input layer (bottom layer) is numbered 1, while the output layer (top layer) is numbered 4. The first example shows the structure that results if a single stimulus is placed in the visual field. The structure of the 'attentional beam' is important here and is shown in Figure 1B. It is this structure that the attentional mechanism attempts to impose on selected stimuli. The second example shows the time course of attentional selection if two stimuli are placed in the visual field. The specifics of network architecture (numbers of units and connections) are arbitrary and variations in these parameters do not affect the conclusions.

In Figure 1A, only the feed-forward connections are shown; the feed-back connections are not shown for simplicity (see Tsotsos et al. 1995b for details on the full circuitry). A stimulus which spans 2 units in the input layer is to be attended by the system; the resulting attentional beam is shown in Figure 1B. The grey lines represent inactive connections, the black lines represent connections whose feedforward flow is inhibited by the attentional beam, and the red lines represent feedforward connections activated by the stimulus. Red units are activated solely by the red stimulus. It is assumed that across the receptive field of each unit, the weights applied to the inputs are Gaussian; otherwise, the signal at the output layer contains no information that can be used by a WTA process to localize the peak. The red shading shows the strength of response.

The WTA mechanism locates the peaks in the response of the output layer of the pyramid, the two remaining red units in Figure 1B. The recursive WTA process then is extended from top to bottom, pruning away the connections that might interfere with the

selected units. Eventually, the two stimulus units are located in the input layer and isolated within the beam.



**Figure 1**

**A. A hypothetical visual processing pyramid showing the portion of the pyramid activated due to the initial feedforward stimulation of a single input stimulus. Grey connections are unaffected by either feedforward or feedback signals. Red connections denote connections along which feedforward signals from the red units flow. In the input layer, the red units represent the input stimulus. White coloured units represent units with zero response. B. The final configuration of the attentional beam reacting to a single input stimulus. Black connections are inhibited connections. Note the 'pass zone' in red that allows the stimulus to reach the output layer, and the 'inhibit zone' in black that prevents any other stimulus from reaching the output layer at the same location as the selected stimulus.**

In Figure 2, the second example, a five-step sequence depicts the changes over time that the visual processing pyramid undergoes if there are two stimuli in the input layer (shown in Figure 2A), using the same network configuration as in the previous figure and again showing only the feed-forward connection. They are colour-coded red and blue as are the connections and units which are activated solely by them (the Gaussian weighting is not shown to simplify the colouring). The mauve coloured units and connections are those which are activated by both stimuli regardless of proportions. In other words, it is assumed that if a unit sees only its preferred stimulus in its receptive field it will respond strongly, and otherwise will respond only weakly. In order then to create the colouring shown, it is further assumed that each of the units in the figure represents a column or assembly of neurons specific for a particular receptive field, and the colour represents the strongest response within that column. Note that much of the pyramid is affected by both stimuli and as a result, most of the output layer gives a confounded response.

The subject is directed to attend to the location of the red stimulus. How is the location cue determined? When giving a subject a location cue for a display, a cue item is displayed at a particular location; the subject's visual system analyzes this spot and notes it is the cue. For the model, a location cue is provided in the same manner. An initial display is provided that is analyzed by the system, the result of which is used as the cue.

The location cue for this example leads to appropriate units corresponding to the location in the output layer are marked as shown in Figure 2B.

The output of those units does not reflect the desired input at this point in time. Rather, these units form the root of the attentional beam as it begins its downward traversal through the visual pyramid. The next phase of the computation is to push the beam down one level further, locating the units which will be the attended ones within the beam. Simultaneously, the feedforward connections from units in layer 3 which feed the attended units in the output layer are inhibited and are not part of the pass zone.

Figure 2C shows the connections in black whose feed-forward flow is explicitly inhibited by the attentional mechanism between layers 3 and 4. Note is that at this early stage of the application of the attentional beam, very little seems to be changing. The large scale changes come later as more of the visual pyramid is affected. The selection of units also moves down one level to layer 2. Figure 2D shows major changes; the next set of connections between the middle layers are inhibited. Those inhibitions cause several units in 3 layer to have no active input and thus they provide no signal to the output layer. Those connections are coloured grey, the same as the other inactive connections in the pyramid. This change in turn causes several units in the output layer previously coded mauve to be coded red; that is, they receive signals originating only from the red input stimulus.

The final stage of the process leads to the network shown in Figure 2E. After this point, the selected units in the output layer receive input only from the selected stimulus in the input layer. Note that several units in the output layer are coded in blue, showing that the effect of the blue stimulus still gets through the beam structure, in fact stronger than in the unattended case of Figure 2A. These signals may be available for processing without awareness.

As well, a few of the weakly responding mauve units still remain. The interference between stimuli evident in Figure 2A is eliminated completely with respect to the item attended, and much reduced for the unattended item. The events depicted with this set of figures would occur in the 100 to 200 ms after stimulus onset (for example, as shown in Chelazzi et al. 1993). Note the difference between the pattern of activations in Figures 1B and 2E. In the former case, no location cue is given; the winner-take-all mechanism chooses the strongest responses in the output layer and inhibits the rest. Thus the set of connections and units attended forms the structure shown with the active connections being strictly those permitted by the selection mechanism. In the latter case, a location cue is given; thus, there is no inhibition within the output layer (if there were, none of the blue or mauve units would survive).




Figure 2

A. The visual processing pyramid at the point where the activation due to two separate stimuli in the input layer has just reached the output layer. No attentional effects are yet in evidence. B. Attention is focused at the location of the output layer corresponding to the location of the selected input. C. The first level of inhibition due to the attentional beam. The feedforward flow of the black connections is inhibited. D. The second level of attentional inhibition. Several units in layer 3 now receive no input and thus do not provide signals to the output layer. E. The third and final level of inhibition due to the attentional beam.

### 3. Observing Neural Attentive Enhancement and Suppression

Motter (1993) describes an elegant experiment where a number of important observations were made. The experiment itself involved an array of stimuli (oriented bars) on which monkeys were trained to select a cued stimulus from the array (3 to 8 stimuli) and report the orientation of the stimulus. The stimuli were presented in a circular array; a fixation target was centered in this circle. Attention was directed either toward the receptive field

that was being recorded or away from the receptive field by the presentation of a location cue on the array prior to the stimulus array. The observations may be summarized as:

- changes in responses of V1, V2 and V4 neurons were observed under certain conditions of focal attention;
- both relative increases and relative decreases in neural activity were observed in association with attention being directed at receptive field stimuli;
- the presence of competing stimuli were a major factor in determining the presence or absence of differential sensitivity. Most neurons were not differentially sensitive with single stimuli;

Is it possible to explain these results with the selective tuning model for the above example? It suffices for the purpose of providing explanations to the above observations that there be at least two stimuli in the input layer and that one is selected as the attended one. Location cues are given by showing the system an image of the cue location (a cross on a blank background for example; see Tsotsos et al. 1995b). The cue image is processed by the same mechanism, it is attended and the beam that attends is thus set up for the stimulus. To show Motter's effects, it is necessary to track the time course of qualitative responses for a subset of units in the network of the example in Figure 2, and some are identified by arrows in Figure 3. Figure 4 summarizes the whole network. It is clear that both kinds of neural changes are expressed during attentive processing.

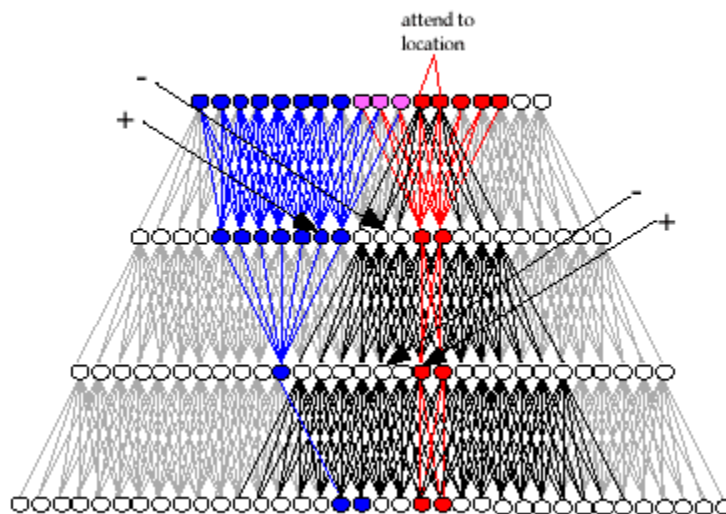


Figure 3

Examples of units within each of the two middle layers of this network are shown where enhancement and suppression of responses in the time course of applying attentive selection are observed. In order to draw the comparison to Motter (1993), it is important that each of these units

have receptive fields that contain the attended location (the red stimuli). Each unit is connected to 7 units in the layers above and below and thus this requirement is met.

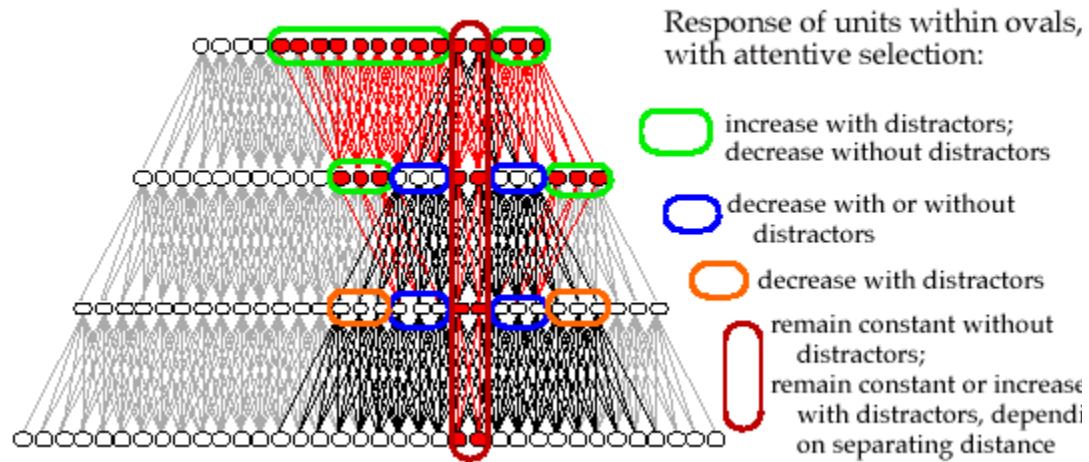


Figure 4

This figure summarizes the changes that occur at each level of the processing network for both the case when distractors are present within the beam as well as when they are not. By "not present", it is meant that they are not present within the attentional beam. It should be clear that if distractors are outside the beam, then this is equivalent to there being no distractors for the attended stimulus.

## 4. Impact on the Triangular Circuit Model

LaBerge describes a model that is not matched to the level of complexity of the system it is trying to explain. From the huge number of connections among visual areas, thalamus and frontal areas, LeBerge has selected subsets of 3 and linked them to his claims with no further enabling mechanisms. Although it is accepted that the connections between the thalamus and virtually all other cortical visual areas exist, it is not so obvious that the only explanation for their function is to provide enhancement for attentively selected neurons. The selective tuning model shows that single-cell observations of neural enhancement or suppression are a result of a more subtle process than LaBerge or others describe. Although the mechanism may be inhibitory at a local level, the effect is both enhancement and suppression due to the structure of the network. Further, LaBerge does not explain how routing of an attended signal might occur. The solution to routing is control of information flow at all layers of the visual processing hierarchy, a feature that is not present in the LaBerge model.

But there is a deeper set of issues, more basic, that LaBerge is side-stepping in setting up the thalamus as an attentional control center (ACC). What are the goals that an ACC must



satisfy? At least 3 tasks seem important: selection of stimuli to attend, routing of stimuli through the processing hierarchy in order to remove interfering contextual stimuli, and coordination in space and time so relevant neurons throughout the hierarchy are all attending to the same items.

What information does the ACC require in order to accomplish these tasks? In order to select the items to attend, it seems necessary to have a global view of the visual field; no algorithm could determine the most salient item otherwise. Task instructions modify the selection of most salient item so that the selection reflects current goals and therefore ACC must have access to those instructions. In order to localize the attended stimulus in retinotopic space as well as in feature space the abstraction accomplished by the feedforward paths must be reversed. In other words, in feedforward processing, precise location information associated with features is discarded (many-to-one neural mapping). The reversal may occur by using a local search process in the feedback direction to determine what is the most important or salient item in a neuron's input. Such a search process must occur for every neuron that exhibits a feedforward many-to-one mapping and must be done in such a manner so that all relevant neurons attend to the same stimulus. The information required for this is local to each neuron. The mapping at each level of the processing hierarchy is different; how can the thalamus compute all the relevant transformations?

What is the connectivity required for ACC to receive this information? Since it cannot be determined in advance where salient stimuli appear in the visual field, full connectivity from each neuron in each layer of the processing hierarchy seems necessary in order to provide a path for the information to the ACC.

What processing must ACC perform in order to satisfy its goals? In order to detect the most salient item in the visual field, it has been shown that a winner-take-all model operating on a representation of saliency will suffice (Koch & Ullman 1985; Tsotsos et al., 1995). Although the question of what form this representation takes is not yet fully answered, nor is it known whether a biological correlate to WTA exists, it does appear that a centralized saliency representation would require far greater connectivity than is observed in any visual area while a distributed representation of saliency does not exhibit this problem. A centralized ACC faces the same problems that a single centralized representation of saliency does because it subsumes such a representation.

What is the connectivity required for ACC to communicate its decisions to the visual processing hierarchy? Due to the dual requirements of modulation of the pathway that an attending stimulus takes through the hierarchy and ensuring that all neurons along the path are attending to the same stimulus, ACC must have feedback connections to each neuron in each layer of the processing hierarchy.

The above set of requirements form a rather daunting task for any central structure, such as the thalamus, to accomplish. LaBerge must demonstrate how the required computations can be performed in neural response times and how connectivity needs are met.

## References

Anderson, C., & Van Essen, D. (1987). Shifter circuits: A computational strategy for dynamic aspects of visual processing. *Proc. Nat. Acad. of Science USA*, 84, 6297-6301.

Chelazzi, L., Miller, E., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, 363, 345-347.

Connor, E.C., Preddie, D., Gallant, J., & Van Essen, D. (1997). Spatial attention effects in macaque area V4. *J. Neuroscience*, 17(9), 3201-3214.

Hernandez-Peon, R., Scherrer, H., & Jouvett, M. (1956). Modification of electrical activity in the cochlear nucleus during attention in unanesthetized cat. *Science*, 123, 331-332.

LaBerge, D. (1997). Attention, Awareness and the Triangular Circuit. *Cognition and Consciousness*, 6, 149-181.

Motter, B. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2 and V4 in the presence of competing stimuli. *Journal of Neurophysiology*, 70(3), 909-919.

Nyberg, L., McIntosh, A., Cabeza, R., Nilsson, L.-G., Houle, S., Habib, R., & Tulving, E. (1996). Network analysis of positron emission tomography regional cerebral blood flow data: Ensemble inhibition during episodic memory retrieval. *Journal of Neuroscience*, 16(11), 3753-3759.

Tsotsos, J.K. (1995a). Towards a Computational Model of Visual Attention. In T. Pappas, C. Chubb, A. Gorea, & E. Kowler (Eds.), *Early vision and beyond*. (pp. 207-218). Cambridge, MA: MIT Press.

Tsotsos, J.K., Culhane, S., Wai, W., Lai, Y., Davis, N., & Nuflo, F. (1995b). Modeling visual attention via selective tuning. *Artificial Intelligence*, 78, 507-547.

Tsotsos, J.K. (1990). Analyzing Vision at the Complexity Level. *Behavioral and Brain Sciences*, 13(3), 423-445.

Tsotsos, J.K. (1993). An Inhibitory Beam for Attentional Selection. In L. Harris & M. Jenkin (Eds.), *Spatial vision in humans and robots*. (pp. 313-331). Cambridge: Cambridge University Press.

Van Essen, D., Anderson, C., Felleman, D. (1992). Information Processing in the Primate Visual System: An Integrated Systems Perspective. *Science*, 255(5043), 419-422.

von der Malsburg, C. (1981). The correlation theory of brain function. (Internal Report 81-2). Göttingen, FRG: Dept. Neurobiology, Max-Planck-Institute for Biophysical Chemistry.