

Varieties of Attention: A Model of Visual Search

Giedrius T. Buračas, Thomas D. Albright, and Terrence J. Sejnowski

*Systems Neurobiology Lab, Salk Institute,
10010 N. Torrey Pines Rd., La Jolla, CA 92037*

ABSTRACT

We have trained monkeys to perform a feature conjunction search task for color and motion and have recorded from neurons in area MT during the performance of this task. In order to put the experimental results into a theoretical context, we have developed a system-level model of visual processing incorporating several attentional mechanisms known to function in mammalian visual systems. A reinforcement learning (temporal difference) algorithm was employed to replicate the learning process in monkeys. The model learned to perform the feature-conjunction search task with performance closely resembling that of human and monkey conjunction search.

The model builds on the notion of two visual streams: The temporal visual stream, crucial for object recognition, exerts top-down influences on early visual representations; these influences (feature-specific attention) prime feature detectors to bias their sensitivity towards object features of interest (feature selection). The parietal (dorsal) visual stream, known to be involved predominantly in spatial vision (coordinate transformations for various actions), exerts top-down spatial selection on the feature maps. Both feature and spatial selection processes bias bottom-up activation of feature maps so that both stimulus salience and the behavioral goals are reflected in the resulting *saliency map* (Koch & Ullman, 1985) which determines the site of information readout into memory.

INTRODUCTION

Despite the central place held by attention in cognitive psychology, research on attention has been marked by contradictions between competing schools of thought (e.g. early vs. late filtering; negative definition of attention as processing bottleneck vs. attention as a positive adaptation for selection of information for behavioral ends; etc.). In this paper we suggest how various manifestations of visual attention might be explained by a distributed system of visual processing. We attempt to build a bridge between visual search behavior, known to involve dynamic allocation of attention, and physiological data relevant to the operation of attention. The bridge includes three levels of analysis (Marr, 1982): (i) computational goals of attention, (ii) algorithms of selection by attention, and (iii) the mechanisms of selection and their impact on sensory information processing. In the following sections we summarize psychophysical and physiological data obtained from monkeys performing a visual search task, and present a systems-level model to account for the findings.

Two broad processing streams have been identified in the visual system: processing for action in the dorsal visual stream (transformations of spatial information), and processing for object recognition and identification in the ventral visual stream (shape and feature analysis; e.g., Milner & Goodale, 1995). We suggest that massive back-projections in these two streams of processing may mediate two different top-down control mechanisms, performing spatial and featural selection of targets for further processing.

The serial search task, on which our experiments and modeling is focused, is likely to engage both spatial selection and feature selection mechanisms. We propose that, on one hand, feature selection biases feature representations (in the spatially global way) in early visual cortical maps so that (i) regions gaining positive bias become interpreted as "figure" (vs. non-selected "ground"); (ii) the effective search-set gets constrained (c.f. Wolfe et al, 1989; Wolfe, 1994; Desimone & Duncan, 1995; Duncan, 1996). On the other hand, focused spatial selection (i) supplies ascendancy to the selected regions of cortical "feature" maps so that these regions gain priority for recognition and awareness over nonselected regions (cf. Koch & Ullman, 1985); (ii) modifies neuronal population codes to represent finer information about local stimulus properties. Thus, the two selection mechanisms interact to mutually constrain processing to those aspects of visual stimuli most relevant to behavioral goals.

In the following sections we first describe our psychophysical data obtained from humans and monkeys, and results of single unit recording from area MT of macaques performing a serial visual search task. Then we present a system level model for visual search and detailed simulations of attentional modulations in feature maps. These simulations attempt to account for (sometimes contradictory) neurophysiological data on attentional modulations. In closing, we show that the systems-level model can be trained to perform the visual search task.

EXPERIMENTAL RESULTS

PSYCHOPHYSICS

We adopted a paradigm of visual search for conjunctions of color and motion task, which requires voluntary attention (Nakayama & Silverman, 1986). We trained human subjects and a rhesus monkey on the search task diagrammed in Figure 1.

As shown previously (Nakayama & Silverman, 1986), the search time for a unique combination of color and direction of motion in multi-object arrays increases linearly with the number of distractors for human observers. Search times collected in psychophysical experiments with monkeys, trained to perform the color/motion conjunction search, also depend linearly on number of distractors (Fig.2 a,b). Despite significant decrease (up to six-fold) in the slope of the search function of some human subjects with practice, the slopes remain comparable to that of monkeys (6.7-23.1 ms/aperture for humans, 7.5-17.2 ms/aperture, for monkeys) (Buracas & Albright, 1994). The monotonically increasing search times provide evidence that both species perform the task in a manner demanding




Fig. 1. The search task. A sample aperture --the cue -- is presented initially for one sec., followed by a search array of apertures containing random dots moving to the left or right and colored in one of the two colors (red or green). Subject's task is to determine whether the sample conjunction of color and direction of motion is present in the array. Humans had to press one key if the target is present, and another key if absent. Monkeys were required to make a saccade to the target. RF indicates the locus of a receptive field. The arrows, indicating the direction of motion inside apertures are not shown during experiments.

serial allocation of attention, which recruits posterior parietal cortices in humans (as suggested by PET data; Corbetta et al, 1995).

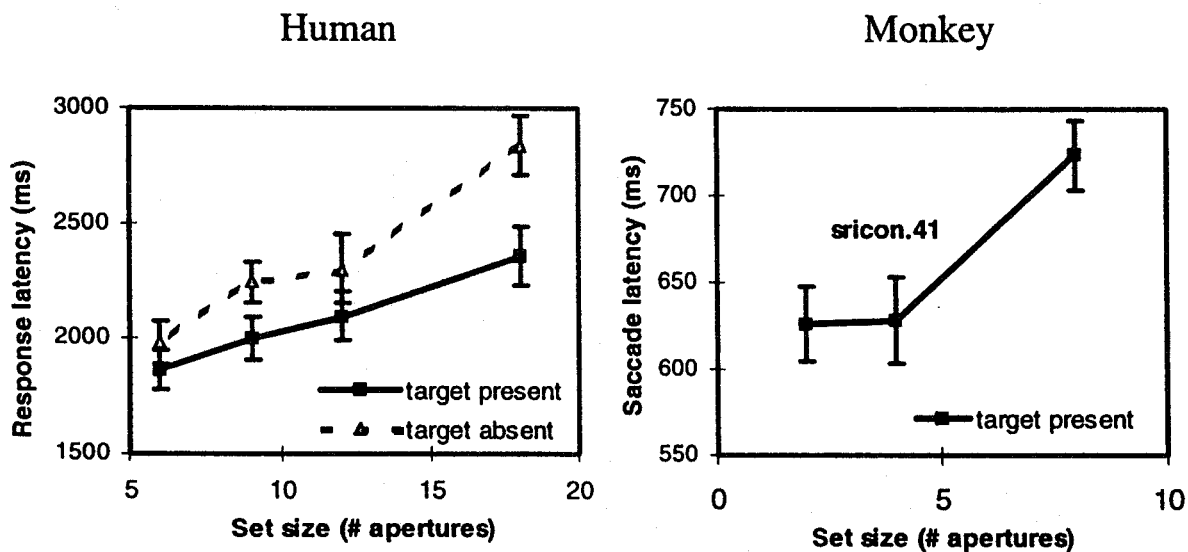


Fig. 2. Searching for conjunction of color and motion by a human (left) and a monkey (right). Both search slopes for target present (TP) and target absent (TA) conditions are shown for humans, and TP only shown for a monkey. The intercepts for humans and monkeys vary because of differences in behavioral response (humans press a button, monkeys make a saccade to the target).

NEUROPHYSIOLOGY

We recently reported that the responses of many neurons in area MT are modulated by attention during a period of tentative target detection by a monkey performing the search task (Fig. 3) (Buracas & Albright, 1995). Invariably, all neurons significantly affected by modulation increased their firing rates during the period of target detection (a small fraction of neurons decreased their firing rates when attention was *withdrawn* from the RF). This new evidence that area MT neurons are strongly affected by attention lays the foundation for our hypothesis that MT neurons, like neurons in other "feature maps" (e.g. V4, V2), are influenced by top-down selection mechanisms.

The effect of attention on discriminability of target's direction of motion is revealed by the receiver operating characteristic (ROC) for an individual neuron: our analysis shows that an ideal observer, by using attentional modulation, can substantially increase its chance of correct inference of target's direction (e.g., from 68% to 93% correct for a representative neuron). This result suggests that attention might increase *information rates* about stimulus associated with MT spike trains.

In conclusion, our neurophysiological experiments have revealed a highly consistent and novel attentional modulation in area MT (cf. Treue & Maunsell, 1995) which may be interpreted as the effect of attention gating information rates in area MT.

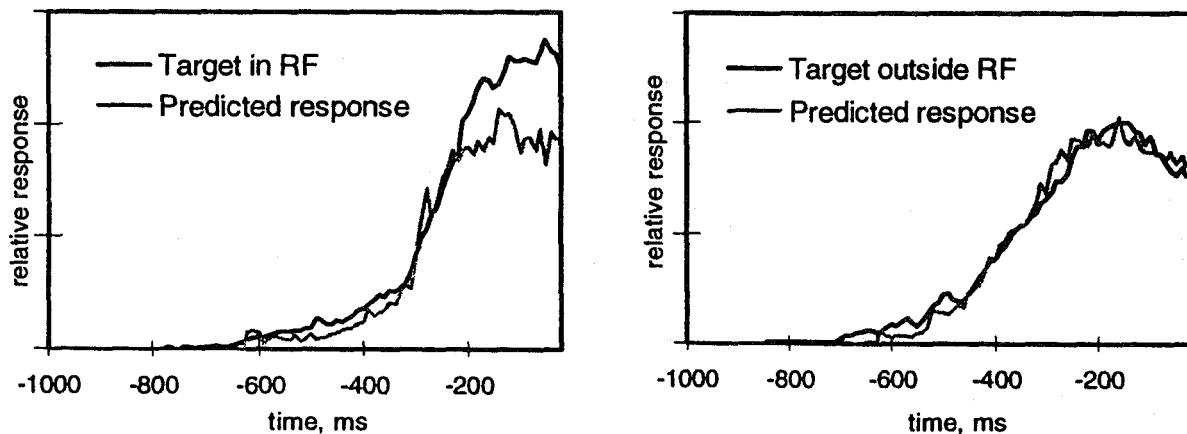


Fig. 3. Saccade-aligned histograms with corresponding estimated baselines (gray line -- responses, estimated from target-absent trials) for target in receptive field condition (left) and target outside receptive field condition (right). When target is inside RF and the target is correctly detected, the neuron, which encoded the motion of the target, invariably increased its firing rate by an average of 30% in about one third of MT neurons. No such response facilitation was observed for target outside receptive field (right). The estimated baseline (gray lines) was calculated from neuronal responses during target absent trials during which monkeys were fixating eyes for one second.

MODELING

These experimental results set the stage for a model whose goal is to explain (i) how various brain regions interact to produce the search behavior; and (ii) identify the neuronal mechanisms responsible for the attentional modulation observed by us in area MT and reported by other laboratories.

THE SYSTEMS-LEVEL MODEL OF VISUAL SEARCH: THE ARCHITECTURE

The temporal sequence of neuronal events in the model starts with activation of feature maps and proceeds up to decision stages as follows: First, the search array consisting of a set of apertures containing a colored moving texture inside them is presented, and activates neurons selective for color and motion in feature maps (Fig.4). Each of the two feature maps contains two "retinotopic" submaps for two different feature values (8x8 unit array with torus topology, which prevents undesirable edge effects). The neuronal representations of stimuli is highly simplified at this point: neurons are assumed to be feature detectors, responding only to their preferred feature values, i.e. a neuron's response I_{ij} is defined

$$I_{ij}(u) = R_{ij} \exp(-(u - u_j)^2 / 2\sigma_{ij}^2), \quad (1)$$

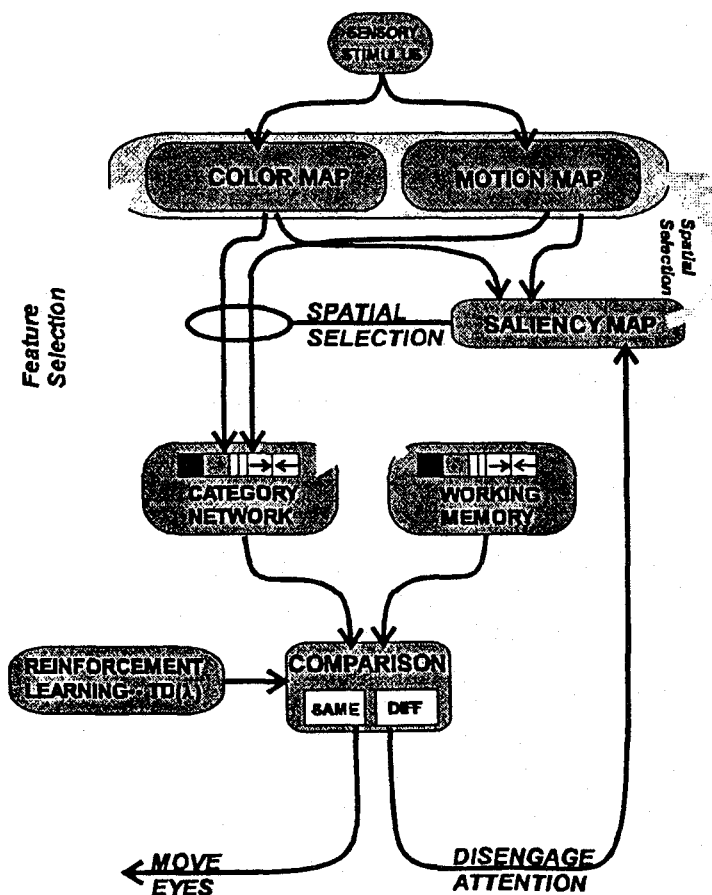


Fig. 4. The flow of neuronal events in the large-scale model of visual search. See text for details.

with stimulus values u satisfying either $u = u_j$, or $|u - u_j| > 3\sigma_{ij}$, where σ_{ij} determines the width of a tuning curve, $i = [x, y]$ is a position vector, and $j \in [1..N]$ indexes preferred feature values (a more elaborate treatment will be presented in the next subsection).

It is well established that receptive fields of many neurons in visual cortical feature maps are surrounded by inhibitory surrounds which could enhance feature contrast (e.g., V1 -- Knierim & Van Essen, 1992; MT -- Miezin et al., 1985; V4 -- Desimone et al., 1993). Thus type of contrast enhancement is introduced by convolving each feature submap with a kernel L of lateral inhibition:

$$x_{ij} = \max\left(0, \int L(a)I(i-a)da\right), \quad (2)$$

where the ratio of positive/negative weight volume for the kernel $L(i)$ is 2:1 (thus, the kernel is not balanced; cf. Wolfe, 1994). The N feature detectors from one map (e.g. red and green detectors within a color map) with overlapping receptive fields then undergo further competition as discussed at length in the following section.

ATTENTIONAL MODULATIONS IN FEATURE MAPS

Neuronal responses in primate extrastriate visual areas can undergo dramatic modulation in animals performing tasks that demand allocation of attention. It is well established that neurons in parietal, inferior temporal cortices, area V4, and, to a lesser degree, V2 and V1, can be affected by attention (Wurtz et al, 1980; Haenny et al, 1988; Moran & Desimone, 1985; Motter, 1993, 1994a,b; Spitzer et al, 1988). New evidence obtained by us and others has added area MT to the list of cortical areas affected by attention (Buracas & Albright, 1995; Treue & Maunsell, 1995). In the following discussion the analysis is constrained to attentional modulation in cortical "feature maps" relevant (to a first approximation) to the color/motion conjunction search task, i.e. areas MT and V4, thought to represent motion and color respectively (e.g. Zeki, 1978; Albright, 1984).

Attentional modulation in the two cortical feature maps appear to share many common features. When a single stimulus is projected into a receptive field (RF) of a V4 neuron, responses to it may significantly increase when the stimulus is used in a behavioral task (Fisher & Boch, 1985; Haeny and Schiller, 1988; Haeny et al, 1988; Motter, 1993). A similar moderate enhancement is observed in area MT when attention is directed to a stimulus in a receptive field (Buracas & Albright, 1995; Treue & Maunsell, 1995). Furthermore, if an animal is subject to a difficult orientation or color discrimination task, tuning curves of V4 neurons acquired during this task appear narrower and taller when compared to tuning curves from an easy discrimination task (Spitzer et al, 1988).

A somewhat different picture emerges when two stimuli are projected into a receptive field of a neuron in a feature map. In general, if the two stimuli possess antagonistic features along some dimension, the neuronal response to the composite stimulus usually is less than to a single optimal stimulus. For example, in V1, which has an orientation map, if an orthogonal bar is overlaid on top of a bar eliciting an optimal response, the response will decrease -- a phenomenon, known as cross-orientation

inhibition. Similarly, two random-dot fields moving in preferred and null direction inside a RF of an MT (motion map) neuron elicits weaker response than a single random dot field, moving in the neuron's preferred direction (Snowden et al., 1991). Preliminary data from area V4 also suggests a similar antagonistic mechanism (Reynolds et al, 1995). These data suggest a universal mechanism operating in cortical feature maps -- local competitive circuits. When attention is applied to a receptive field containing two stimuli, the modulatory effect seems to be consistent with suppressing the interfering effect of the second, non-optimal, stimulus (Moran & Desimone, 1985; Reynolds et al, 1995). Similar effects were observed in area MT with stimuli moving inside a receptive field in opposite directions (Treue & Maunsell, 1995). Thus, it seems likely that attentional modulation can affect not just individual neurons but the outcome of local competitive interactions as well.

The attentional modulation discussed above are spatially localized, and thus are controlled by some top-down mechanism of spatial selection. In addition, human data (e.g. Anllo-Vento & Hillyard, 1996; Buracas & Albright, 1996; Egeth et al, 1984;) and neurophysiological evidence (Motter, 1994a,b) points to the existence of a spatially global but feature-specific mechanism of selection. Most convincingly, Motter (1994a,b) has reported that neurons in V4 can be modulated by the relevance of stimulus color without specifying the location for spatial allocation of attention. These findings suggest that responses of V4 neurons, tuned to a behaviorally relevant stimulus value, can be biased by a feature-specific top-down influence. Interestingly, we did not find such a feature specific biasing (color- or direction-specific) in area MT (Buracas & Albright, 1995).

The data presented above serve as the foundation for a model of attentional modulation in feature maps (Fig.5). A dynamic network model for competition in feature maps is given by:

$$\tau \frac{dy_{ij}}{dt} = G_j c x_{ij}^{a_i} - y_{ij} \left(c + \sum_{k \in P} G_k \cdot (w x_{ik})^{a_i} \right). \quad (3)$$

where inputs $x_{ij}(t)$ compete for representation by the outputs $y_{ij}(t)$ by means of shunting inhibition, the indices i and j refer to spatial location and preferred feature value respectively, the exponent a_i determines the strength of competition controlled by the degree of spatial focusing (correlated with task difficulty) at a site i , the feature specific




Fig. 5. A hypothetical soft competition mechanism in feature maps: Inputs, before reaching the output units are affected by nonlinearity. This nonlinearity might be implemented as a postsynaptic membrane property of the summing and output units (e.g. voltage gated Ca^{++} channels). Note, that the recurrent influence from output y_i is local and can, again, be implemented by shunting ion channels (like Cl channels).




Fig. 6. Temporal evolution of solutions to the competition network, Eq. (3). Horizontal axis represents preferred feature values, vertical axis plots responses of a set of neurons with overlapping receptive fields. The input profile portrays a response of a population of neurons, tuned to different values, when a single stimulus is presented. The feature bias profile portrays top-down feature-specific gating strength as a function of neurons preferred feature, G_i . Right: Evolution of responses for two different values of a . G_i is set to 1 for all i . Left: The effect of feature-specific bias to the outcome of competition.

and spatially global factors G_j bias the response in one or another direction, and the constants c and w allow to define how the competition affects the height and width of tuning curves. Summation is over the set P of all neurons tuned to different values at a single retinotopic location. Typical behavior of the (1) is presented in Fig.6. The steady-state solution of Eq. (3) is given by:

$$y_{ij} = \frac{G_j c x_{ij}^{a_i}}{c + \sum_{k \in P} G_k (w x_{ik})^{a_i}} \quad (4)$$

For a fixed Gaussian input pattern, as a_i increases from 1 to 1.5 the tuning curve undergoes a change matching closely to simple scaling, but with larger a_i the tuning curve grows and narrows until its height reaches the saturation point, but the width keeps shrinking (Fig.7). This kind of behavior accounts for apparently contradictory reports on tuning curves in area V4: Relatively easy tasks only scale tuning curves (e.g., Motter, 1993), while a hard discrimination task, besides upscaling can also significantly narrow the tuning curve.

These two aspects of spatial attentional modulation -- scaling and narrowing -- may be a reflection of its two different computational goals: (i) Spatial selection of objects in a visual scene for further processing in the ventral stream. Simple boosting of relevant neuronal responses would suffice for this purpose, since higher tiers of visual processing in




Fig. 7. The effect of attentional focus on tuning curves. Left: with increasing a a Gaussian tuning curve grows first and then starts narrowing. Right: comparison of a tuning curve at value $a=1.5$ with a scaled version of the original tuning curve. The subtle difference between scaling and increasing a is hard to detect experimentally for $1 < a < 1.5$.

this stream possess competitive mechanisms not unlike the one discussed here (Chellazi et al, 1993). (ii) Refining representations of stimuli for scrutiny or fine discrimination, since attention can facilitate performance of discrimination tasks (e.g. Palmer, 1990). Narrowing of tuning curves is an efficient way of increasing the resolution power of a neuronal population code.

Eq. (4) can also account for data obtained with competing stimuli inside a receptive field: If a nonpreferred stimulus is added to the preferred stimulus in the receptive field when a_i is set to 1 (attention removed), the response of a neuron will decrease significantly. Increasing the value of a_i (i.e. the degree of spatial focusing) will bring the response back to the original level, as observed in MT and V4 (Treue & Maunsell, 1995; Reynolds et al., 1995).

In conclusion, a broad variety of attentional modulation affecting neuronal responses in cortical maps necessitate multiple selection mechanisms subserving different behavioral ends.

The feature maps are affected by spatial and feature selection mechanisms (cf. Nowlan & Sejnowski, 1995). Based on electrophysiological evidence on modulatory effects of both spatially-focused and feature-based attention in visual cortical areas, we propose that spatial attention (or attentive effort) controls the strength of local competition, but feature-based attention serves to bias feature representations in direction of values currently stored in the working memory (Fig.4). We elaborate on the mechanism of spatial selection in the next section.

SALIENCY MAP

In the second step of processing, responses in different feature maps project into a common topographic saliency map (Koch & Ullman, 1985), implemented as weighted averaging:

$$s_i(t+1) = \sum_{j=1..4} w_j y_{ij}(t) - IR_i(t) + n(t), \quad (5)$$

where $w_j=0.25$ in our simulations, j ranges over all feature values, and $n(t)$ is zero-mean Gaussian noise; the discrete time t refers to the ordinal number of a search cycle. In addition there is an adjunct topographic "inhibition of return" map which keeps track of visited locations by activating units, pointing to locations of current maximum saliency. This map controls disengagement of attention. Visited locations have a short-term memory, determined by the decay factor γ : The inhibition-of-return map is updated at every time step, i.e. $IR_i(t+1) = \gamma \cdot IR_i(t)$, with $\gamma < 1$. Thus, attention is shifted to the second highest saliency peak by setting an inhibition-of-return map's unit, which points to the current peak of saliency, to $IR_i=1$. This procedure is similar to that, originally proposed by Koch & Ullman (1985), and used, among others, by Wolfe (1994) and Pouget & Sejnowski (1996).

In the current version of our model, the maximum activation of the saliency map determines the spatial locus in feature maps, from which feature values are fed into the nontopographic "category network", containing four units that correspond to abstracted representations of 4 possible feature values. At any given time, the category units represent feature values of an object to which the spatially-focused attention is pointing. An alternative feature readout architecture is suggested by data of Chelazzi et al. (1993), indicating that the final target selection by a competitive mechanism may occur in the category network (i.e. inferior temporal cortex).

The feature values activated in the category network then can be compared with the target feature values in working memory stored while viewing the cue (see caption of Fig.1 for a description of the behavioral paradigm). The results of comparison determine further course of action: If the features in working memory match the activated features in the category network, the response "target found" is executed (e.g., eyes are moved to the target). In case of a mismatch, attention is disengaged by resetting the saliency map as described above. The process is repeated until the target is found.

Computing the comparison of working memory and category network unit activations is at heart of the serial search task, since it is the outcome of comparison which drives the behavior of the system. Thus, it is likely that learning the correct mapping from results of comparison to actions is the essential component of learning the feature conjunction search task for monkeys. We modeled monkey-like learning behavior by adding a module that learned the mapping from working memory-category network comparison results (match/no match) to action (move eye/ move attention) solely from a reinforcement signal. Since the serial search is an inherently sequential behavior, the




Fig. 8. Learning to search for conjunctions of features by a model (30 training sessions/point).

temporal difference learning algorithm (Sutton, 1988) was a suitable method for learning this task. The learning module (Fig. 4, see also Fig. 10) consisted of an evaluation network that learned to predict reward and a policy network that learned to map working memory-category network comparison results onto actions (Barto et al, 1983; Sutton, 1988; Montague et al, 1995).

The model of visual search described above processes visual information in the search array in a selective and serial fashion. Feature-specific bias, propagating from working memory constrains the region of potential allocations of spatial attention, while the spatially-focused attention determines sites of detailed representation of feature values in the category network (note the symmetry of mutual interactions in the two attentional systems).




Fig. 9. Search functions generated by the trained model. The search functions are plotted for several values of feature specific feedback. The numbers on the right are gain factors G_i , used to boost responses of feature detectors corresponding to target feature values (cf. Wolfe, 1994).

THE SEARCH BEHAVIOR OF THE MODEL

The systems-level model outlined above is capable of learning to search for conjunctions of features. In this section we present some results on learning performance and model's behavior during visual search for conjunctions of features.

We trained the model to perform the conjunction search task by means of reinforcement learning algorithm (e.g., Montague et al, 1995). Prior to every search trial a new random search array was generated and the target feature conjunction was stored in the working memory module. Then, for every search cycle the sequence of computations was performed as outlined at the beginning of Modeling section. If the Policy Network (Fig.10) made a correct decision to "move eyes" when the S-map ("attentional focus") was pointing to the target's location, the positive reward was assigned, and the negative reward was generated if the "move eyes" command was issued while the attentional pointer was pointing to a distractor. Zero reward was assigned after "move attention" command. The network learned to perform the search task within 500 trials (Fig.8).

The trained model was tested on a variety of conjunctive and disjunctive search tasks. The search behavior exhibited by the model matched closely the performance of monkeys and humans (Buracas & Albright, 1996): while disjunctive (feature) search induced search times independent of the search-array size, the search for conjunctions of features demands serial examining of the search array, thus, causing search times to increase monotonically with the number of distractors (Fig. 9). Increasing the feature-

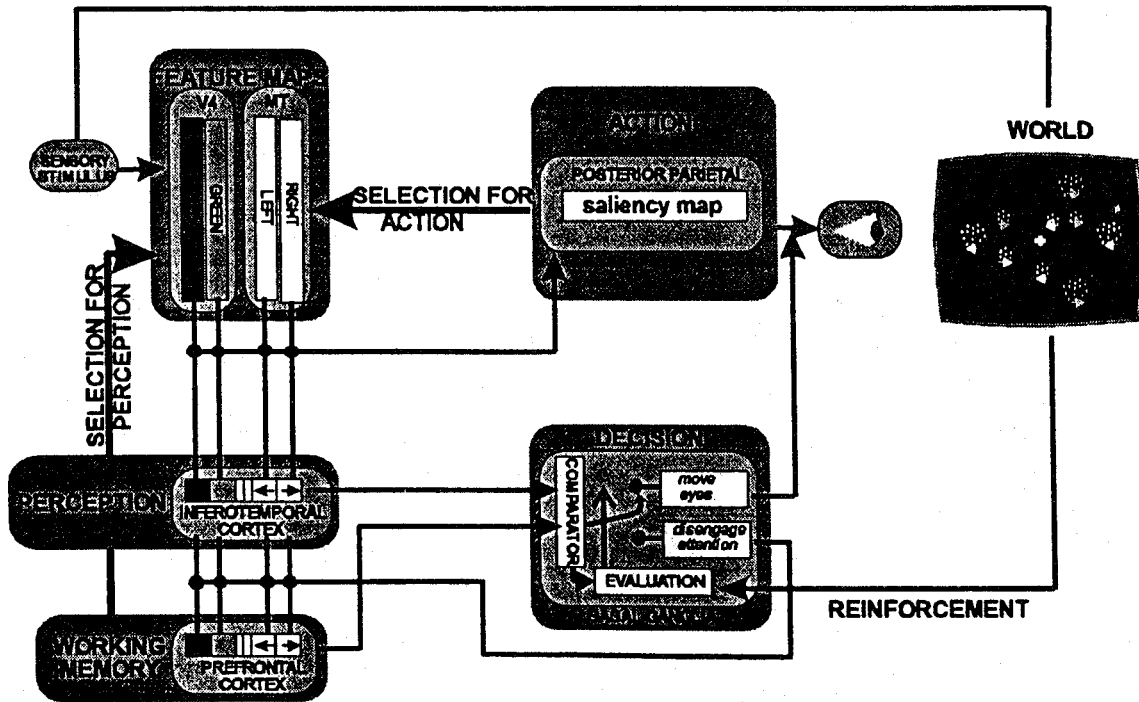


Fig. 10. Mapping components of the Visual Search Model onto the brain regions.

REFERENCES

- Albright, T.D. (1984) *J Neurophysiol*, **52**, 1106-1130.
- Anllo-Vento, L. & Hillyard, S.A. (1996) *Perception & Psychophysics*, **58**, 191-206.
- Barto, A.G., Sutton R.S., Anderson, C.W. (1983) *IEEE Transactions on Systems, Man, and Cybernetics*, **13**, 834-846.
- Berns, G.S. & Sejnowski, T.J. (1996) In: *Neurobiology of Decision-Making*. Springer-Verlag.
- Buračas, G.T. & Albright, T.D. (1994) *Soc. Neurosci. Abs.* **20-1**, 318.
- Buračas G.T. & Albright, T.D. (1995) *Soc. Neurosci. Abs.* **21-3**, 1759.
- Buračas G.T. & Albright, T.D. (1996) *Invest. Ophthalm. & Visual Sci.*, **37**, s298.
- Chelazzi, L., Miller, E.K., Duncan, J.U., & Desimone, R. (1993) *Nature*, **363**, 345-347.
- Corbetta, M., Shulman, G., Miezin, F.M., and Petersen, S.E. (1995) *Science*, **270**, 802-805.
- Desimone, R. & Duncan, J. (1995) In: *Annual Review of Neuroscience*, **18**, 193-222.
- Desimone, R., Moran, J., Schein, S.J., and Mishkin, M. (1993) *Visual Neurosci.*, **10**, 159-171.
- Duncan, J. (1996) In: *Attention and Performance*, **XVI**, 550-578.
- Farah, M.J. (1990) *Visual Agnosia*. MIT Press.
- Fisher, B., and Boch, R. (1985) *Brain Res*, **345**, 111-113.
- Goldman-Rakic, P.S., & Selemon, L.D. (1990) *TINS*, **13**, 241-244.
- Haenny, PE, and Schiller, PH. (1988) *Exp Brain Res*, **69**, 245-259.
- Haenny, PE, Maunsell, JHR, and Schiller, PH. (1988) *Exp Brain Res*, **69**, 225-244.
- Knierim, J.J. & Van Essen, D.C. (1992) *J. Neurophysiol.*, **67**, 961-980.
- Koch, C, and Ullman, S (1985) Shifts in selective visual attention: towards the underlying neural circuitry. In: Vaina, LM (1987) *Matters of Intelligence*, D.Reidel Publishing Company, 115-141.
- Marr, D (1982) *Vision*. W.H.Freeman and Co. New York.
- Miezin, F., McGuinness, E., and Allman, J. (1985) *Annual Review of Neuroscience*, **8**, 407-30.
- Miller, E.K., Erickson, C.A., and Desimone, R. (1995) *Soc. Neurosci. Abs.*, **21**, 661.
- Milner, A.D. and Goodale, M.A. (1995) *The Visual Brain in Action*. Oxford University Press.
- Montague, P.R., Dayan, P., Person, C., and Sejnowski, T.J. (1995) *Nature*, **377**, 725-728.
- Moran, J, & Desimone, R (1985) *Science*, **229**, 782-784.
- Motter, B (1993) *J. Neurophysiol*, **70**, 909-919.
- Motter B (1994a) *J Neurosci*, **14**, 2178-2189.
- Motter, B (1994b) *J Neurosci*, **14**, 2190-2199.
- Nakayama K, Silverman GH (1986) *Nature*, **320**, 264-265.
- Nowlan, S.J., and Sejnowski, T.J. (1995) *J. Neurosci.*, **15**, 1195-2919.
- Palmer, J. (1990) *J. Exper. Psychol.: Human Perception and Performance.*, **16**, 332-350.
- Rafal, R.D. & Posner M.I. (1987) *Proc. Natl. Acad. Sci. U.S.A*, **84**, 7349-53.

- Reynolds, J., Nicholas, J., Chelazzi, L., Desimone, R. (1995) *Soc. Neurosci. Abs.* **21-3**, 1759.
- Selemon, L.D., & Goldman-Rakic, P.S. (1985) *J. Neurosci.*, **5**, 776-794.
- Snowden, R.J., Treue, S., Erickson, R.G, and Andersen, R.A. (1991) *J. Neurosci.*, **11**, 2768-2785.
- Spitzer, H, Desimone, R, and Moran, J (1988) *Science*, **240**, 338-340.
- Sutton, R.S. (1988) *Machine Learning*, **3**, 9-44.
- Treue, S. & Maunsell, J.H.R., (1995) *Soc. Neurosci. Abs.* **21-3**, 1759.
- Treisman, AM & Gelade, G (1980) *Cog. Psychol.*, **12**, 97-136.
- Wolfe, JM, Cave KR, & Franzel, SL (1989) *J.Exper.Psychol.*, **15**, 419-433.
- Wolfe, JM (1994) *Psychonomic Bulletin & Review*, **1**, 202-238.
- Wurtz, R. H., Goldberg, M. E., and Robinson, D. L. (1980) *Prog. Psychobiol. Physiolog. Psychol.*, **9**, 43-83.
- Zeki, S.M. (1978) *J. Physiol. (London)*, **277**, 273-290.