

Reprinted from

PROGRESS IN BRAIN RESEARCH

VOLUME 121

DISORDERS OF BRAIN, BEHAVIOR AND COGNITION: THE NEUROCOMPUTATIONAL PERSPECTIVE

EDITED BY

JAMES A. REGGIA

*Departments of Computer Science and Neurology, UMIACS,
University of Maryland,
College Park, MD 20742, USA*

EYTAN RUPPIN

*Departments of Physiology and Computer Science,
Schools of Medicine & Mathematics,
Tel-Aviv University, Tel-Aviv 69978, Israel*

DENNIS GLANZMAN

*Theoretical and Computational Neuroscience Research Program,
The Neuroscience Center, National Institute of Mental Health, NIH,
Bethesda, MD 20892, USA*



ELSEVIER

AMSTERDAM - LAUSANNE - NEW YORK - OXFORD - SHANNON - SINGAPORE - TOKYO

1999

Frames of reference in hemineglect: a computational approach

A. Pouget^{1,*}, S. Deneve¹ and T.J. Sejnowski²

¹*Brain and Cognitive Science Department, University of Rochester, Rochester, NY, 14627, USA*

²*Howard Hughes Medical Institute, The Salk Institute for Biological Studies, La Jolla, CA 92037, USA and Department of Biology University of California, San Diego La Jolla, CA 92093, USA*

Unilateral lesions of the parietal lobe in humans often lead to a syndrome known as hemineglect (Heilman et al., 1985; Pouget and Driver, 1999). These patients experience difficulty processing or reacting to stimuli located in the hemispace contralateral to their lesion. For instance, after a right parietal lesion, the case we consider here, patients fail to eat food located on the left side of their

plates and forget to dress on the left side of their bodies. When asked to cross out line segments displayed on a page, a clinical test known as line cancellation, they fail to cross out the lines located on the left side of the page (Fig. 1).

Although it is clear that hemineglect patients have problems with the left side of space, the frame of reference of this left hemispace has remained difficult to delineate. In principle, these patients could be neglecting only objects falling on their right retinal hemifield, in which case hemineglect

*Corresponding author. e-mail: alex@bcs.rochester.edu

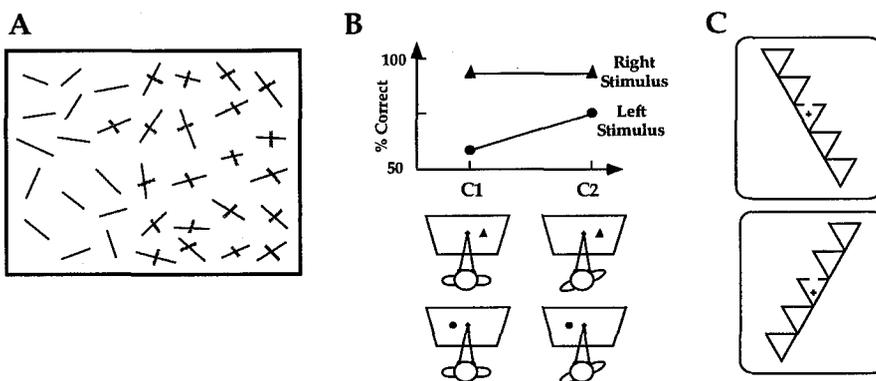


Fig. 1. A. Typical result for a line cancellation test in a left neglect patient. B. Percentage of correct identification in the Karnath et al. experiment (1993). In condition 1 (C1), subjects were seated with eyes, head, and trunk lined up whereas in condition 2 (C2) the trunk was rotated by 15° to the left. The overall pattern of performance is not consistent with pure retinal or pure head-centered neglect and suggests a deficit affecting a mixture of these two frames of reference. C. The two displays used in Driver et al. (1994) experiment. Patients must detect a gap in the upper part of the central triangle. In the top (resp. bottom) display, the object made out of the triangles is perceived as rotated 60 degrees clockwise (resp. counterclockwise). Left parietal patients detect the gap more reliably in the bottom display, i.e., when the gap is associated with the right side of the object.

would be purely retinocentric. Alternatively, they could be neglecting objects located on the left of their trunk, which would correspond to trunk-centered neglect. Other frames of reference such as head-centered or environment centered (allocentric) are also possible.

These various possibilities can be disentangled by fixing the position of the stimulus in one frame of reference while testing the subject in multiple postures. An experiment by Karnath et al. (1991) provides a good example of this strategy. Subjects were asked to identify a stimulus that could appear on either side of the fixation point. In order to test whether the position of the stimuli with respect to the body affects performance, two conditions were tested: a control condition with head straight ahead (C1) and a second condition with head rotated 15° on the right (where right is defined with respect to the trunk) or, equivalently, with the trunk rotated 15° to the left (where left is defined with respect to the head) (see Fig. 1B, C2). In C2, both stimuli occurred further to the right of the trunk than in C1, though at the same location with respect to the head and retina. Moreover, the trunk-centered position of the left stimulus in C2 was the same as the trunk-centered position of the right stimulus in C1.

As expected, subjects with right parietal lesions performed better on the right stimulus in the control condition (C1), a result consistent with both retinotopic and trunk-centered neglect. However, to distinguish between the two frames of reference, performance should be compared across conditions. If the deficit was purely retinocentric, the results should be identical in both conditions since the retinotopic locations of the stimuli do not vary. On the other hand, if the deficit is purely trunk-centered, the performance on the left stimulus should improve when the head is turned right since the stimulus now appears further toward the right of the trunk-centered hemispace. Furthermore, performance on the right stimulus in the control condition should be the same as performance on the left stimulus in the rotated condition since they share the same trunk-centered position in both cases.

Neither of these hypotheses is fully consistent with the data. As expected from retinotopic neglect, subjects always performed better on the right

stimulus in both conditions. However, performance on the left stimulus improved when the head was turned right (C2), though not sufficiently to match the level of performance on the right stimulus in the control condition (C1, Fig. 1B). Therefore, these results suggest a retinotopically based form of neglect modulated by trunk-centered factors. In addition, Karnath et al. (1991) tested patients on a similar experiment in which subjects were asked to generate a saccade toward the target. The analysis of reaction time revealed the same type of results as the one found in the identification task, thereby demonstrating that the spatial deficit is, to a first approximation, independent of the task. Several other experiments have found that neglect affects a mixture of frames of reference in a variety of tasks (Bisiach et al., 1985; Calvanio et al., 1987; Ladavas, 1987; Ladavas et al., 1989; Farah et al., 1990; Behrmann & Moscovitch, 1994).

At first glance, these results would appear to be consistent with current views regarding spatial representations in the parietal cortex. Several authors argue that the parietal cortex contains a mosaic of cortical areas each encoding the location of an object in *one* particular frame of reference and each involved in the control of *one* particular behavior (Goldberg, 1990b; Stein, 1992; Colby, 1998).

Large lesions, which are unfortunately quite common, would be expected to encompass many of these areas and the resulting neglect would affect multiple frames of reference. However, neglect would also be expected to be behavior specific. For instance, it has been suggested that cells in LIP use retinotopic coordinates for the control of saccadic eye movements (Goldberg et al., 1990) whereas cells in VIP use head-centered receptive fields for the control of reaching toward the face (Graziano et al., 1994; Duhamel et al., 1997). A lesion of these modules would affect both behaviors but the deficit would be retinotopic for eye movements and head-centered for reaching. More generally, for any particular behavior, the deficit should be confined to one particular set of coordinates. The fact that Karnath et al. (1993) found that a mixture of frame of reference was affected for different tasks, verbal report vs. saccades, is incompatible with this prediction.

In this chapter, we suggest a different explanation for the mixture of frames of reference found in hemineglect subjects. Our approach is based on interpreting the responses of parietal neurons as basis function of their sensory inputs. We show that basis function decomposition is a computationally efficient way to represent the position of objects in space for the purpose of sensorimotor transformation. Basis function representations also have the property that each cell can encode the location of objects in *multiple* frames of reference simultaneously allowing them to be involved in the control of multiple behaviors.

Thus, we concur with previous reports that the parietal cortex appears to contain multiple modules, but we argue that, in all modules, each neuron encodes *multiple* frames of reference. As a result, a lesion leads to a syndrome in which each behavior is affected in multiple frames of reference in all patients.

In the final section, we also address the issue of object-centered neglect, which refers to the possibility that patients neglect the left side of object regardless of the position and orientation of objects in space. Object-centered neglect is clearly illustrated in an experiment by Driver et al. (1994) in which patients were asked to detect a gap in the upper part of a triangle embedded within a larger object (Fig. 1C). They reported that patients detected the gap more reliably when it was associated with the right side of the object than when it belonged to the left side even when this gap appeared at the same retinal location across conditions (Fig. 1C).

We show that the basis function framework can also account for such results without invoking the existence of explicit object-centered representations. We also discuss other claims about object-centered neglect and argue that, in many cases, the deficit is not necessarily in object-centered coordinates but rather in a relative viewer-centered coordinate frame (the term viewer centered here means any frame of reference centered on a part of the body, such as the trunk or the eyes). We call this relative neglect and we demonstrate that basis functions combined with a temporal selection process are sufficient to account for this behavior.

This chapter is organized as follow: It starts with a brief summary of the basis function approach followed by the description of the architecture used for the simulations of hemineglect. Next, the behavior of a lesioned basis function network is examined in a line cancellation task. Then the issue of frame of reference is addressed by simulating the Karnath experiment. Finally, object-centered neglect is examined in the light of the basis function approach.

Basis function representation

The basis function model of parietal cortex is motivated by the hypothesis that spatial representations correspond to a recoding of the sensory inputs that facilitates the computation of motor commands. This perspective is consistent with Goodale and Milner's (1990) suggestion that the dorsal pathway of the visual cortex mediates object manipulation (the 'How' pathway), as opposed to simply localizing objects as Mishkin et al. (1983) previously suggested (the 'Where' pathway). In general, the choice of a representation strongly constrains whether a particular computation is easy to perform or difficult. For example, addition of numbers is easy in decimal notation but difficult with Roman numerals. The same is true for spatial representations. With some representations the motor commands for grasping may be simple to perform and stable to small input errors, but in others the computation could be long and sensitive to input errors.

A set of basis functions has the property that any *nonlinear* function can be approximated by a *linear* combination of the basis functions (Poggio, 1990; Poggio and Girosi, 1990). Therefore, basis functions reduce the computation of *nonlinear* mappings to *linear* transformations – a simpler computation. Most sensorimotor transformations are *nonlinear* mappings of the sensory and posture signals into motor coordinates; hence, given a set of basis functions, the motor command can be obtained by a *linear* combination of these functions. In other words, if parietal neurons compute basis functions of their inputs, they would recode the information in a format that simplifies the computation of subsequent motor commands.

As illustrated in Fig. 2B, the response of parietal neurons can be described as the product of a Gaussian function of retinal location multiplied by a sigmoid function of eye position. Sets of both Gaussians and sigmoids are basis functions, and the set of all products of these two basis functions also forms basis functions over the joint space (Pouget and Sejnowski, 1995, 1997b). These data are therefore consistent with the assumption that parietal neurons compute basis functions of their inputs and thus provide a representation of the sensory inputs from which motor commands can be computed by simple linear combinations (Pouget & Sejnowski, 1995, 1997b).

It is important to emphasize that not all models of parietal cells have the properties of simplifying

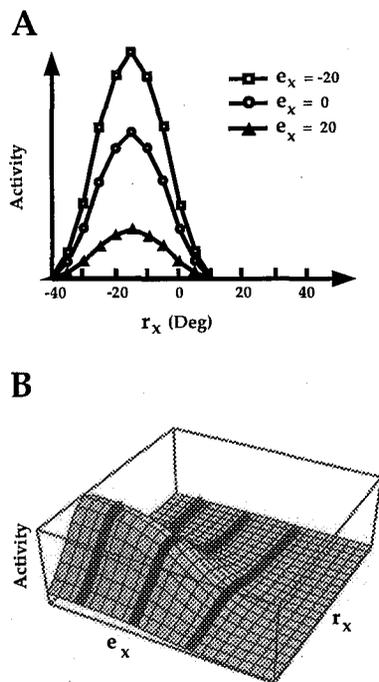


Fig. 2. A. Idealization of a retinotopic visual receptive field of a typical parietal neuron for three different gaze angles (e_x). Note that eye position modulates the amplitude of the response but does not affect the retinotopic position of the receptive field (adapted from Andersen et al., 1985). B. 3D plot showing the response function of an idealized parietal neuron for all possible eye and retinotopic positions, e_x and r_x . The plot in A was obtained by mapping the visual receptive field of this idealized parietal neuron for three different eye positions as indicated by the bold lines

the computation of motor commands. For example, Goodman and Andersen (1990) as well as Mazzoni and Andersen (1995) have proposed that parietal cells simply add the retinal and eye position signals. The output of this linear model does not reduce the computation of motor commands to linear combinations because linear units cannot provide a basis set. By contrast, the hidden units of the Zipser and Andersen model (1988), or the multiplicative units used by Salinas and Abbott (1995, 1996), have response properties closer to the basis function units, and the basis function hypothesis can be seen as a formalization of these models (for a detailed discussion see Pouget and Sejnowski, 1997b).

One interesting property of basis functions, particularly in the context of hemineglect, is that they represent the positions of objects in multiple frames of reference simultaneously. Thus, one can recover simultaneously the position of an object in retinocentric *and* head-centered coordinates from the response of a group of basis function units similar to the one shown in Fig. 2B (Pouget and Sejnowski, 1995, 1997). As shown in the next section, this property allows the same set of units to perform multiple spatial transformations in parallel.

This approach can be extended to other sensory and posture signals and to other parts of the brain where similar gain modulations have been reported (Trotter et al., 1992; Boussaoud et al., 1993; Bremmer and Hoffmann 1993; Field and Olson 1994; Brotchie et al., 1995). When generalized to other posture signals, such as neck muscle proprioception or vestibular inputs, the resulting representation encodes simultaneously the retinal, head-centered, body-centered, and world-centered coordinates of objects.

We recently explored the effects of a unilateral lesion of a basis function network (Pouget and Sejnowski, 1997a). The next section describes the structure of this model.

Model organization

The model contains two distinct parts: a network for performing sensorimotor transformations and a

selection mechanism. The selection mechanism is used when there is more than one object present in the visual field at the same time.

Network architecture

The network has three layers with basis function units in the intermediate layer to perform a transformation from a visual retinotopic map input to two motor maps in head-centered and oculo-centric coordinates respectively (Fig. 3). The visual inputs correspond to the cells found in the early stages of visual processing and the set of units encoding eye position have properties similar to the neurons found in the intralaminar nucleus of the thalamus (Schlag-Rey and Schlag 1984). These input units project to a set of intermediate units that contribute to both output transformations. Each intermediate unit computes a Gaussian of the retinal location of the object, r_x , multiplied by a sigmoid of eye position, e_x :

$$O_{ij} = \frac{e^{-(r_x - r_{xi})^2/2\sigma^2}}{1 + e^{-\beta(e_x - e_{xj})}} \quad (1)$$

We consider horizontal positions only because the vertical axis is irrelevant for hemineglect. These units are organized in two two-dimensional maps covering all possible combinations of retinal and eye position selectivities. The only difference between the two maps is the sign of the parameter β , which controls whether the units increase or decrease activity with eye position.¹ β was set to 8° for one map and -8° for the other map. The indices (i, j) refer to the position of the units on the maps. Each location is characterized by a position for the peak of the retinal receptive field, r_{xi} , and the midpoint of the sigmoid of eye position, e_{xj} . These quantities are systematically varied along the two dimensions of the maps in such a way that in the upper right corner r_{xi} and e_{xj} correspond to right retinal and right eye positions whereas in the lower left they correspond to left retinal and left eye positions.

This type of basis function is consistent with the responses of single parietal neurons found in area 7a. A population of such units forms a basis function map encoding the locations of objects in head-centered and retinotopic coordinates simultaneously. The activities of the units in the output maps are computed by a simple linear combination of the basis function unit activities. Appropriate values of the weights were found by using linear regression to achieve the least mean square error.

This architecture mimics the pattern of projections of the parietal area 7a, which innervate both the superior colliculus and the premotor cortex (via the ventral parietal area VIP; Andersen et al., 1990a; Colby and Duhamel, 1993), where neurons have retinotopic and head-centered visual receptive fields respectively (Sparks, 1991; Graziano and Gross, 1994).

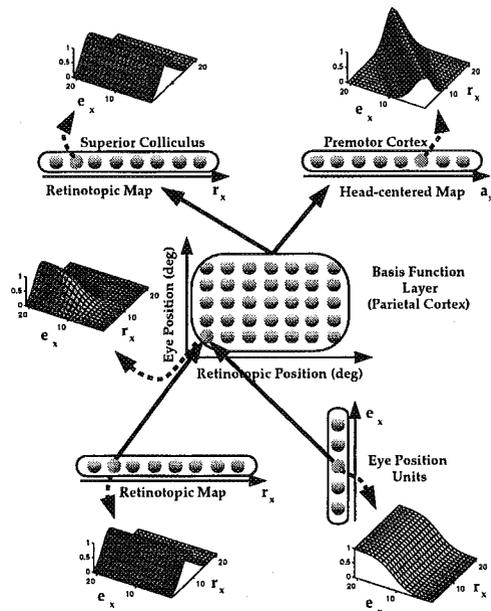


Fig. 3. Network architecture. Each unit in the intermediate layers is a basis function unit with a Gaussian retinal receptive field modulated by a sigmoid function of eye position. This type of modulation is characteristic of the response of parietal neurons.

¹ We use the term *retinal* position to refer to the position of the object in space with respect to a frame of reference centered on the eye. This is not to be confused with the position of the object onto the retina. Indeed, due to the optics of the eye, the image of an object in the *right* retinal hemifield will be projected onto the *left* hemiretina.

Hemispheric biases and lesion model

Although the parietal cortices in both hemispheres contain neurons with all possible combinations of retinal and eye position selectivities, most cells tend to have their retinal receptive fields on the contralateral side (Andersen et al., 1990a,b). Whether a similar contralateral bias exists for the eye position in the parietal cortex remains to be determined although several authors have reported such bias for eye position selectivities in other parts of the brain (Schlag-Rey and Schlag, 1984; Galletti and Battaglini, 1989; Van Opstal et al., 1995).

In the model, we divide the two basis function maps into two sets of two maps, one set for each hemisphere (again the two maps in each hemisphere correspond to two possible values for the parameter β). Units are distributed across each hemisphere to create neuronal gradients. These neuronal gradients induce more activity overall in the left maps than in the right maps when an object appears in the right retinal hemifield *or* when the eyes are turned to the right, with the opposite being true in the right maps.

Several types of neuronal gradients can lead to such activity gradients. The gradients we used for the simulations presented here affected only the maps with positive β ; that is, maps with units whose activity increased when the eyes turn right. In both the right and left maps, the number of units for a given pair of (r_{xi}, e_{xi}) increased linearly with contralateral values of eye and retinal location, as indicated in Fig. 4. The slope of the gradient was chosen to be steeper for retinal location than for eye position. This choice was motivated by electrophysiological recordings. As mentioned above, the retinal gradient is clearly present in experimental data whereas the eye position gradient is not as clear. This would therefore suggest that if the eye position gradient exists – and our model predicts that it does – it is weaker than the retinal gradient.

A right parietal lesion was modeled by removing the right parietal maps and studying the network behavior produced by the left maps alone. The effect of the lesion was therefore to induce a neuronal gradient such that there was more activity in the network for right visual field and right eye positions.

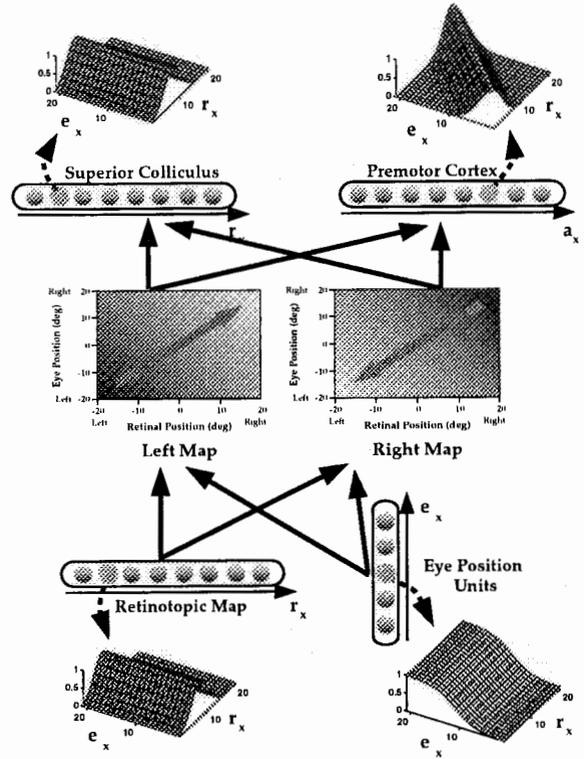


Fig. 4. Same as in Fig. 3 but with the intermediate layer split into right and left hemisphere maps. The arrows indicate the direction of the neuronal gradient. Each map overrepresents contralateral retinal and eye positions. A right parietal lesion is simulated by removing the right map.

The exact profile of the neuronal gradient across the basis function maps did not matter as long as it induced a monotonically increasing activity gradient as objects were moved further to the right retinal position and the eyes fixated further to the right. The results presented in this chapter were obtained with linear neuronal gradients.

Selection model

The selection mechanism in the model was adapted from Burgess (1995) and was inspired by the visual search theory of Treisman and Gelade (1980) and the saliency map mechanism proposed by Koch and Ullman (1985). Unfortunately, most of these existing models are not based on distributed representations of the kind used in the present model. Models of stimulus selection, for instance, typically use local representations in which a

stimulus is characterized by one number, usually the activity of a single units (Burgess, 1995). In contrast, distributed patterns of activity occur in the basis function maps of the present model to represent one or several stimuli (see Fig. 7B).

Therefore, we had first to reduce the dimensionality of our representation before an existing model of target selections could be used. A saliency value was assigned to each stimulus present on the display. The saliency was defined as the sum of the activity of all the basis function units whose receptive field were centered exactly on the retinal position of the stimulus, r_i . This method is mathematically equivalent to defining the saliency of the stimulus as the peak in the profile of activity in the superior colliculus output map.

Once the saliencies of the stimuli were computed, the selection mechanism was initiated. It involved a repetition over time of three steps: winner-take-all, inhibition of return and recovery. On the first time step, the stimulus with the highest saliency was selected by winner-take-all, and its corresponding saliency was set to zero to implement inhibition of return. At the next time step, the second highest stimulus was selected and inhibited, while the previously selected item was allowed to recover slowly. These operations were repeated for the duration of the trial. This procedure ensured that the most salient items were not selected twice in a row. However, due to the recovery process, the stimuli with the highest saliencies might be selected again if displayed for a sufficiently long time.

In this model of selection, the probability of selecting an item is proportional to two factors: the absolute saliency associated with the item and the saliency relative to the ones of competing items.

It is possible to implement a selection mechanism similar to the one described here by using lateral connections within the basis function map (Cohen et al., 1994). This implementation does not make an artificial distinction between the representation and the selection mechanism, as made here, and is more biologically plausible. We favor the idea that there is no distinction between the cells responsible for encoding the location of objects and the ones responsible for selective attention. However, the model of Cohen et al.

(1994) required complex dynamics and computation-intensive simulations and would have produced the same results as the present model. Thus, the selection mechanism used here was motivated by practical considerations.

Evaluating network performance

We used this model to simulate several experiments in which patient performance was evaluated according to reaction time or per cent of correct response.

In reaction time experiments, we assumed that processing involves two sequential steps: target selection and target processing. Target selection time was assumed to be proportional to the number of iterations, n , required by the selection network to select the stimulus using the mechanism described above. Each iteration was arbitrarily chosen to be 50 ms long. This duration matters only when more than one stimulus is present, so that distractors could delay the detection of the target by winning the competition.

The time for target processing (that is to say, target recognition, target naming, etc.) was assumed to be inversely proportional to stimulus saliency, s_i . Thus, the total reaction (RT) time was given by:

$$RT = 100 + 50n + \frac{500}{1000s_i}. \quad (2)$$

The percentage of correct responses to a stimulus was determined by a sigmoid function of the stimulus saliency:

$$p = \frac{0.5}{(1 + \exp(-(s_i - s_0)/t))} + 0.5, \quad (3)$$

where s_0 and t are constants.

This model for evaluating the output is based on signal detection theory when signal and noise follow Gaussian distributions of equal variance (Green and Swets, 1966). This is equivalent to assuming that the rate of correct detection (hit rate) is the integral of the probability distribution of the signal from the decision threshold to infinity.

In line bisection experiments, subjects were asked to judge the midpoint of a line segment. In the network model, the midpoint, m , was estimated

by computing the center of mass of the activity induced by the line in the basis function map.

$$m = \frac{\sum_{\text{all units}} a_i r_{xi}}{\sum_{\text{all units}} a_i}, \quad (4)$$

where r_{xi} is the retinal position of the peak of the visual receptive field of unit i .

All the results given here were obtained from the lesioned model in which the right basis function maps had been removed. For control tasks on the normal network see Pouget and Sejnowski (1997a).

Line cancellation

We first measured the saliency of a stimulus as a function of its retinal location ($[-30^\circ, 30^\circ]$) and eye position ($[-30^\circ, 30^\circ]$) after a simulated right parietal lesion in the basis function network. Saliency was defined as the summed activity in the basis function layer in response to the stimulus.

As shown in Fig. 5A, the saliency increased as the stimulus was moved to the right hemifield *or* when the eyes fixated toward the right. This saliency gradient is a direct reflection of the underlying neuronal gradient introduced by the lesion. The fact that the saliency is affected by both

the retinal location of stimuli and the position of the eyes can account for the fact that neglect affects multiple frames of reference in the model, as elaborated in the next sections.

Next, we tested the network on the line cancellation test, in which patients are asked to cross out short line segments uniformly spread over a page. To simulate this test, we presented the display shown in Fig. 5B and ran the selection mechanism to determine which lines get selected by the network. As illustrated in Fig. 5B, the network only crossed out the lines located in the right half of the display, mimicking the behavior of left neglect patients in the same task (Heilman et al., 1985). The rightward gradient introduced by the lesion makes the right lines more salient than the left lines. As a result, the rightmost lines always won the competition, preventing the network from selecting the left lines. The probability that the line was crossed out as a function of its position in the display is shown in Fig. 5C, where position is defined with respect to the frame of the display. There was a sharp jump in the probability function such that lines on the right of this break had a probability near 1 of being selected whereas lines on the left of the break had a probability near zero (Fig. 5C).

The sharp jump in the probability of selection stands in contrast to the smooth and monotonic

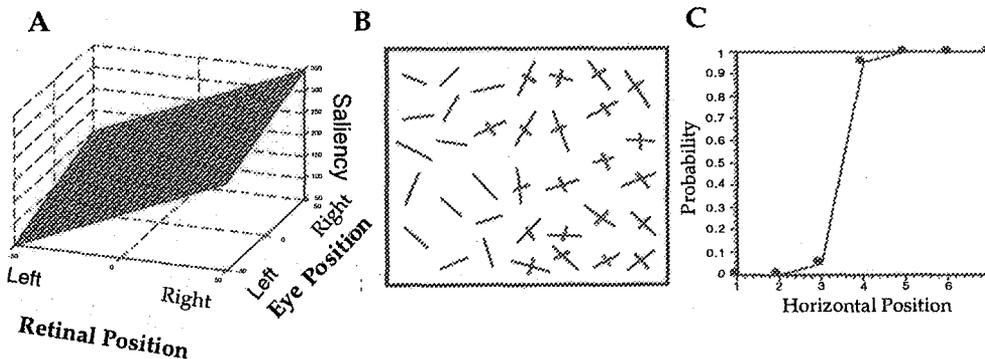


Fig. 5. A. Saliency gradient in a network with a right lesion. Saliency increases when an object is moved toward the right hemiretina or when the eyes move right. B. Simulation of a line cancellation task. The network failed to cross out the line segments on the left side of the page, as in right parietal patients. C. Probability of crossing a line as a function of its horizontal position in the display. The lesioned network fails to cross out the lines in the left half of the display as if the neuronal gradient introduced by the lesion was a step function. The gradient, however, is smooth, and the sudden change in behavior in the center of the display is the result of the dynamics of the selection mechanism.

profile of the neuronal gradient. Whereas the sharp boundary in the pattern of line crossing may suggest that the model ‘sees’ only one half of the display, the linear profile of the neuronal gradient shows that this is not the case. The sharp jump is mainly a consequence of the dynamics of the selection process: because right bars are associated with higher saliencies, they consistently win the competition to the detriment of left bars. Consequently, the network starts by selecting the bar furthest on the right and due to inhibition of return moves its way toward the left. Eventually, however, previously inhibited items recover and again win the competition, preventing the network from selecting the leftmost bars. The point at which the network stops selecting bars toward the left depends on the exact recovery rate and the total number of items displayed.

The pattern of line crossing by the network is not the result of a deficit in the selection mechanism, but rather is the result of a selection mechanism operating on a lesioned spatial representation. The network had difficulty detecting stimuli on the left side of space not because it was unable to orient toward that side of space – it would orient to the left if only one stimulus were presented in the left

hemifield – but because the bias in the representation favored the rightmost bars in the competition.

Frames of reference

The frame of reference of neglect in the model was examined next. Since Karnath et al. (1993) manipulated head position, their experiment was simulated by using a basis function map integrating visual inputs with head position rather than eye position. As in patients, the performance of the network was intermediate between retinocentric and trunk-centered neglect (as shown in Fig. 1B). This result is the direct consequence of the profile of the saliency gradient induced by the lesion (Fig. 5A). In Fig. 6, the left and right stimuli are shown in both conditions superposed on the saliency gradient. In both conditions, the right is more salient than the left stimulus and consequently, the performance on the right stimuli is better. In addition, the saliency of the left stimulus increases from condition 1 to condition 2 which account for the better performance on the left when the head is turned right with respect to the trunk.

Note however, that the saliency of the left stimulus in the second condition does not match the saliency of the right stimulus in the first condition

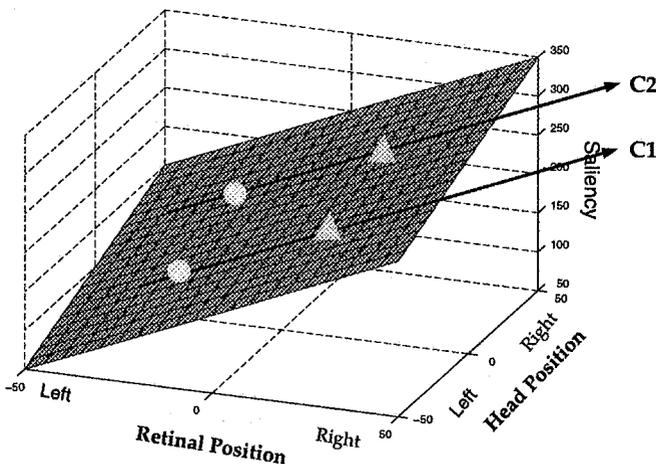


Fig. 6. Right (triangle) and left (circle) stimuli in conditions 1 and 2 of Karnath et al. experiment (Fig. 1B) superposed on top of the saliency gradient. The lesioned model performed better on the right stimulus than left stimulus because the right stimulus was the most salient in both conditions. Additionally, performance for the left stimulus increased from condition 1 to condition 2 due to an increase in saliency. However, the saliency of the left stimulus in condition 2 did not match the saliency of the right stimulus in condition 1 even though they share the same trunk-centered location. This is because the slope of the gradient was steeper along the retinal axis than along the head position axis. As a result, the model did not perform as well on the left stimulus in condition 2 than on the right stimulus in condition 1.

even though they share the same trunk-centered location. This is because the gradient along the retinal axis and the head position axis were not equal. This particular choice of gradient was motivated by physiological constraints as explained in the section on network architecture.

Our simulations were performed with a complete lesion of the right parietal map but, in reality, lesions are incomplete. In some patients, only part of the right hemisphere might be lesioned. The saliency gradient will then reflect the combination of the representation in the left hemisphere and whatever is left in the right hemisphere. It is possible that, in some cases, this gradient would end up being identical along both dimensions in which case neglect would be purely trunk-centered. We would predict that such cases, however, to be very rare.

Therefore, as in humans, neglect in the model was neither retinocentric nor trunk-centered alone but both at the same time. Similar principles can be used to account for the behavior of patients in many other experiments that involve frames of reference (Bisiach et al., 1985; Calvanio et al., 1987; Ladavas 1987; Ladavas et al., 1989; Farah et al., 1990; Behrmann and Moscovitch, 1994).

Object-centered neglect

One of the first studies reporting evidence for object-centered neglect came from the work of Caramazza and Hillis (1990) who reported a patient with a right *word-centered* neglect. This patient made spelling mistakes, primarily on the right side of words, whether the words were presented horizontally, vertically or mirror reversed (in the latter case, the right side of the word appeared in the left hemisphere). Since then, several other studies have observed similar behavior not just for words but for objects as well.

These studies are reviewed in the first part of this section. We argue that in most cases, their results can be explained with what we called *relative neglect*, a form of neglect that does not involve a lesion of an object-centered representation. There are, however, a few studies that cannot be explained by relative neglect such as the one by Driver et al. (1994) described in the introduction. These cases are addressed in the second part of this section.

Recent single cell data suggest a natural extension of the basis function framework to object-centered representations and we explain how this extension of the basis function hypothesis can account for Driver et al. results.

Object-centered neglect or relative neglect?

The paper by Arguin and Bub (1993) provides a good illustration of the kind of data typically used to support object-centered neglect. As shown in Fig. 7A, they found that reaction times were faster when a target (the 'x' in Fig. 7B) appeared on the right of a set of distractors (C2) instead of on the left side (C1), even though the target was at the same retinotopic location in both conditions.

Although this result is certainly consistent with object-centered neglect, it is just as consistent with the idea that patients tend to neglect the parts of the object the furthest to the left, where left is defined with respect to the viewer, *not* the object. In other words, what matters may be the relative position of stimuli, or subparts of an object, along the left-right defined with respect to the viewer. This is what we call *relative neglect*. The result of several other studies can be accounted for with relative rather than object centered neglect. The only way to distinguish between these alternatives is to rotate the object such that the left-right axis of the object is no longer lined up with the left-right axis of the object.

This comment applies as well to standard clinical tests such as line cancellation and line bisection. Line bisection is a test in which patients are asked to judge the midpoint of a line. Left neglect patients typically estimate the midpoint too far to the right. One might be tempted to conclude this rightward overshoot is due to neglect of the left side of the *line* – a line-centered interpretation. If this were the case then rotating the line by 45° with respect to the viewer should have no effect on the performance of the subject. The left side of the line should still be neglected by the same amount. On the other hand, if this is an example of relative neglect, then the overshoot is due to the fact that the subject ignores the part of the line furthest to the left with respect to the subject. This would predict that when the line is vertical, the overshoot should disappear since no part of the line is further to the left than any other

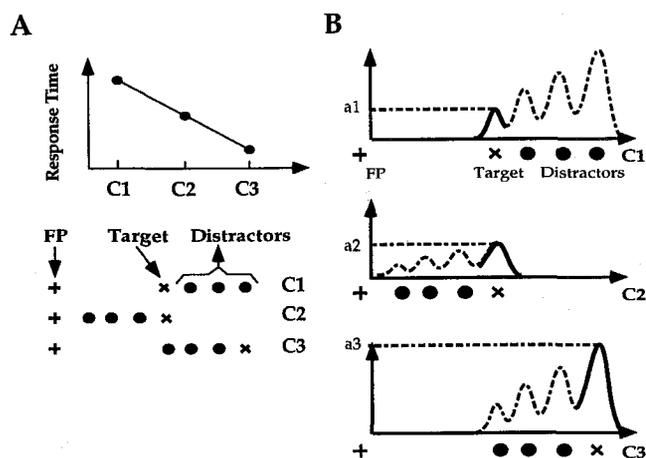


Fig. 7. A. Response times for the Arguin and Bub (1993) experiment for the three experimental conditions illustrated below the graph (FP, fixation point). The decrease from condition 1 (C1) to condition 2 (C2) is consistent with object-centered neglect, i.e., subjects are faster when the target is on the right of the distractors than when it is on the left, even though the retinal position of the target is the same. The further decrease in reaction time in condition 3 (C3) shows that the deficit was also retinotopic. B. Reaction time between conditions 1 and 2 decreased due to the change in the relative saliency of the target with respect to the distractors, even though the absolute saliency of the target was the same in these two conditions ($a_1 = a_2$). FP, Fixation point; C3, condition 3

part. Experimental evidence supports the latter hypothesis. The amount of overshoot is proportional to the cosine of the angle between the line and the viewer with a minimum when the line is near vertical (Halligan and Marshall, 1989; Burnett-Stuart et al., 1991).

Interestingly, relative neglect emerges naturally in our basis function model of neglect. The network reaction times in simulations of the Arguin and Bub (1993) experiments followed the same trends reported in human patients (Fig. 7A). This result can easily be understood if one examines the pattern of activity in the retinotopic output layer of the network for the three conditions in those experiments. Although the absolute levels of activity associated with the target (solid lines) in conditions 1 and 2 were the same, the activity of the distractors (dotted lines) differed in the two conditions. In condition 1, they had relatively higher activity and thereby strongly delayed the detection of the target by the selection mechanism. In condition 2, the distractors were less active than the target and did not delay target processing as much as they did in condition 1.

The network also reproduced the further decrease in reaction time when the whole display, target and distractors, were moved further right on

the retina (condition 3). This indicates that both the relative and absolute retinal position play a role in neglect. This is another example of multiple frames of reference being affected concomitantly in the same patient (and model).

The network also showed relative neglect in a line bisection task. In particular, the overshoot went away for vertical lines (not shown here, but see Pouget and Sejnowski (1997) for more details).

To convincingly establish object-centered neglect, patients should be tested with rotated objects. Object-centered neglect predicts that the neglected part will be the left side of the object whereas relative neglect predicts that it will be the part of the object the furthest on the left with respect to the viewers.

Farah et al. have carried out such an experiment and reported results consistent with relative neglect. Behrmann and Moscovitch (1994) obtained similar results except for asymmetric letters for which the results seemed to indicate object-centered neglect. However, Drain and Reuter-Lorenz (1997) have obtained similar results on asymmetric letters in normals, casting doubt on the object-centered neglect interpretation.

Tipper and Berhmann (1996) have also reported data consistent with object-centered neglect. They

used stimuli made of two circles, one on the left and one on the right. The circles could be linked by a bar, forming a barbell-like object, or not. They explored whether priming can be defined in object-centered coordinates and found that when the right circle was primed, followed by a 180° rotation which brought the right circle to the location of the left circle and vice-versa, the priming stayed with the right circle (now on the left), but *only* if the two circles were linked by the bar.

These results are consistent with the hypothesis that attention can be allocated in object-centered coordinates. A simpler interpretation, however, is that attention is allocated in retinal coordinates and that the attentional spotlight can move with an object. Moreover, the fact that attention followed the left circle only when linked to the other circle would imply that the dynamical aspects of the attentional spotlight are influenced by segmentation factors. Our network model does not include this mechanism and, consequently, cannot account for this experimental result on priming. However, simulations by Mozer shows that Tipper and Behrmann's result can indeed be explained by using an attentional spotlight that works in retinal coordinates (see Mozer's chapter (7) in this book). What makes this demonstration particularly interesting is the fact that Mozer's model was originally designed to perform word recognition. It is therefore possible that Tipper and Behrmann's results are not related to the existence of explicit object-centered representations (a term we define more precisely in the next section) but could be a byproduct of the way the visual system segments and recognizes words and objects.

Ultimately, there appears to be only one result that could not be explained without invoking object-centered representation, namely, the experiment by Driver et al. (1994) described in the introduction and illustrated in Fig. 1C. We show next how this result can be related to the response of single cells using the basis function framework.

Basis function applied to object-centered representations

The existence of object-centered representations at the neuronal level appear to be supported by the

recent work of Olson and Gettner (1995). They trained monkeys to perform saccades to a particular side of an object (right or left, depending on a visual cue) regardless of its position in space. Ideally, it would have been important to test the monkey for multiple orientations of the object but these conditions were omitted in the first study. Once the monkey had acquired the task, they recorded the activity of cells of the supplementary eye field during the execution of the task.

They found that some cells respond selectively prior to eye movements directed to a particular side of an object. For example, a cell might give a strong response before an upward saccade directed to the left side of the object but no response at all to the *same* upward saccade but directed to the right side of the object. This behavior suggests that some cells have motor fields defined in object-centered coordinates which would constitute what we call an explicit object-centered representation. However, *all* the cells recorded by Olson and Gettner can be interpreted as having an oculocentric motor field – a bell-shaped tuning to the direction of the next saccadic eye movement, where direction is defined with respect to the fixation point – which is *gain modulated* by the side of the object and the command (Olson, personal communication).

The gain modulation assumption appears to be supported by recent data from Olson and Gettner, 1998, as well as Breznen et al., 1998. Moreover, this hypothesis is computationally efficient. Performing saccades to the side of an object specified by an instruction can be formalized as a nonlinear mapping from the inputs, the image of a bar and an instruction, to the output, the motor command for the saccade (Deneve and Pouget, 1998). An efficient way to proceed is to use basis functions of the input variables in the intermediate stage of processing. These basis functions could be retinotopic receptive fields gain modulated by the side of the object and the command. Figure 8A shows a basis function neural network that can generate saccades to a particular side of an object according to an instruction and independently of the position and orientation of the object. This is a slightly more general situation than the one considered by Olson and Gettner (1995) since the model can handle arbitrary orientation.

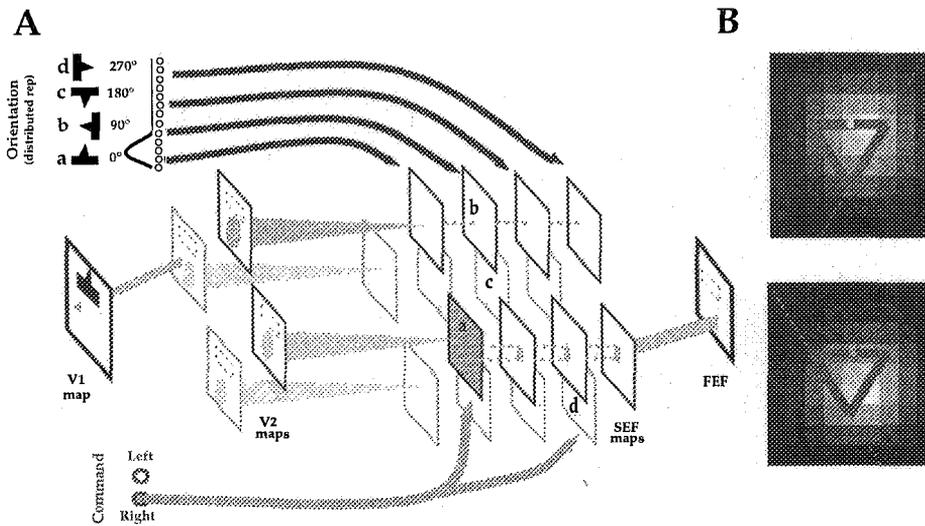


Fig. 8. A. Neural network model for performing saccades to a particular side of an object in response to the image of the object and an instruction. The input contains a V1-like representation of the image of the object, a set of cells tuned to the orientation of the object—similar to cells found in the infero temporal cortex—and a set of cells encoding the current command. The supplementary eye field (SEF) layer contains gain modulated cells like the one found by Olson and Gettner, which compute basis function of the inputs. The motor command in the superior colliculus layer is generated by combining the output of the basis function units. B. Summed activity generated in the SEF in response to the display used by Driver et al. in a lesioned network. The dotted line indicated the general orientation of the object. The activity associated with the upper edge of the triangle is the weakest when the object is tilted clockwise (bottom). This is consistent with Driver's finding that left neglect patients perform worse in this condition.

Deneve and Pouget (1998) have explored the effect of a unilateral lesion in this basis function model. They assumed that the left hemisphere overrepresents right retinal location and counterclockwise object rotation (and vice versa for the right hemisphere). The retinal gradient is identical to the one used by Pouget and Sejnowski (1997) (see the preceding sections). The preference for counterclockwise object rotation is consistent with the general principle used for setting the gradients, namely, the left hemisphere favors any posture toward the right, such as moving the eye to the right, rotating the head to the right, or tilting the head to the right. Indeed, when the head is tilted to the right, the retinal image rotates *counterclockwise*. The hemisphere preferring head tilt to the right should therefore also favor counterclockwise image rotation.

Figure 8B shows the amount of activity in the basis function layer of a lesioned network, in response to the presentation of the triangle display of Driver et al. The important part to focus on is the

upper edge of the central triangle, the edge for which the patients are asked to detect the presence of a gap. As one might expect given the hemispheric gradient, the highest activity is obtained when the object is tilted counterclockwise. This is the case for which subjects perceive the edge to be on the right of the main axis.

This network was not designed to detect the presence of a gap but a simple signal-to-noise argument would predict that the presence of the gap should be more readily detected for counterclockwise than clockwise rotation, as reported by Driver et al. in patients.

The interesting point to note is that the object-centered neglect is obtained here even though the network does not contain an explicit object-centered representation: it does not use cells with motor fields or receptive fields in object-centered coordinates. Instead, the representation uses a more implicit format, involving basis functions of the inputs, which is computationally efficient and consistent with single cell data.

Discussion

The model of the parietal cortex presented here was originally developed by considering the response properties of parietal neurons and the computational constraints inherent in sensorimotor transformations. It was not designed to model neglect, so its ability to account for a wide range of deficits is additional evidence in favor of the basis function hypothesis.

The basis function model captures three essential aspects of the neglect syndrome: (1) it reproduces the pattern of line crossing of parietal patients in line cancellation; (2) the deficit coexists in multiple frames of reference simultaneously; and (3) the model accounts for relative and object-centered neglect. These results rely in part on the existence of monotonic gradients along the retinal and eye position axis (or more generally, the posture axis) of the basis function map. The retinal gradient is supported by recordings from single neurons in the parietal cortex (Andersen et al., 1990), but gradients for the postural signals remain to be demonstrated. The retinal gradient hypothesis is also at the heart of Kinsbourne's theory of hemineglect (Kinsbourne, 1987) and some models of neglect dyslexia and line bisection are based on a similar idea (Mozer and Behrmann, 1990; Anderson, 1996; Mozer et al., 1997).

Other behaviors of hemineglect patients can also be captured by this model, such as the patterns of line bisection or the recovery after vestibular caloric stimulation (see Pouget and Sejnowski, 1997, 1999).

The model presented here cannot account for the fact that left neglect is much more common than right neglect (Heilman, 1985). The reason for this asymmetry is unclear but the common explanation depends on an asymmetry in the hemispheric representations. Whereas the left hemisphere may represent only the right hemifield, the right parietal cortex appears to represent both hemifields in some patients (Kinsbourne, 1987). This would suggest that contrary to what we have assumed in the model, the gradients in the right and left representations are not simply mirror images. Instead, the right hemisphere may have a shallower contralateral gradient and the left hemisphere may have a

steeper gradient. An asymmetric gradient would lead to a preference for the left side of space in a normal network. There is evidence for such a leftward preference in normal subjects and it is therefore possible that the asymmetry observed with right and left neglect is indeed due to a difference in the gradients in the left and right spatial representations (Kinsbourne, 1987; Ladavas, 1990).

One interesting aspect of our approach is that there is no need to represent explicitly all frames of reference to account for the behavior of patients. Instead we have assumed that the position of objects is represented by basis functions, a representation that spans multiple Cartesian frames of reference simultaneously. Consequently, any attempt to determine the Cartesian space in which hemineglect operates is bound to lead to inconclusive results in which Cartesian frames of reference appear to be mixed.

This perspective on frames of reference has interesting implications for rehabilitation of neglect patients. It is often assumed that left neglect patients would improve if one could teach, or force, them to orient toward the left side of space. It has been shown, for example, that neglect improves after caloric stimulation of the vestibular system (consisting of an injection of cold water in the left ear; see Rubens, 1985), or after the presentation of a leftward motion flow field (Pizzamiglio et al., 1990). In both cases, the stimulation induces a series of left eye movements, or nystagmus, forcing the subjects to look toward the left. However, this explanation fails to account for the fact that the recovery persists for several minutes after the left nystagmus stops.

The basis function hypothesis suggests a different interpretation. The neuronal gradient in the lesioned network is such that any change of body posture toward the right improves the saliency of visual stimuli appearing in the left retinal visual hemifield (Fig. 5A). Hence, in the experiment by Karnath et al. (1993; Fig. 1B), the detection and recognition of the left visual stimulus improves when the head is turned toward the right. The same mechanism could explain the effect of the caloric and flow field stimulation. Cold water in the left ear induces an illusion of a head rotation toward the

right. Likewise, a leftward flow field is consistent with a head rotation toward the right. If the brain uses a leaky integrator in both of these situations to compute the amount of rotation, then one would expect the recovery to persist beyond the period of stimulation before slowly disappearing.

It might therefore be worthwhile to pursue rehabilitation methods that take advantage of these general principles; namely, that patients improve when they adjust their posture toward the 'good', or ipsilesional, side of space.

Acknowledgments

This research was supported in part by a fellowship from the McDonnell-Pew Center for Cognitive Neuroscience and an R29 NIH grant to A.P. and grants from the Office of Naval Research and the Howard Hughes Medical Institute to T.J.S. We thank Daphne Bavelier for her comments and suggestions.

References

- Andersen, R.A., Essick, G.K. and Siegel, R.M. (1985) Encoding of spatial location by posterior parietal neurons. *Science*, 230: 456–458.
- Andersen, R.A., Asanuma, C., Essick, G. and Siegel, R.M. (1990a) Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J. Comp. Neurol.*, 296(1): 65–113.
- Andersen, R.A., Bracewell, R.M., Barash, S., Gnadt, J.W. and Fogassi, L. (1990b) Eye position effect on visual memory and saccade-related activity in areas LIP and 7a of macaque. *J. Neurosci.*, 10: 1176–1196.
- Anderson, B. (1996) A mathematical model of line bisection behaviour in neglect. *Brain*, 119: 841–850.
- Arguin, M. and Bub, D.N. (1993) Evidence for an independent stimulus-centered reference frame from a case of visual hemineglect. *Cortex*, 29: 349–357.
- Behrmann, M. and Moscovitch, M. (1994) Object-centered neglect in patients with unilateral neglect: effects of left-right coordinates of objects. *J. Cogn. Neurosci.*, 6(2): 151–155.
- Bisiach, E., Capitani, E. and Porta, E. (1985) Two basic properties of space representation in the brain: evidence from unilateral neglect. *J. Neurol., Neurosurg. Psychiatry*, 48: 141–144.
- Boussaoud, D., Barth, T.M. and Wise, S.P. (1993) Effects of gaze on apparent visual responses of frontal cortex neurons. *Exp. Brain Res.*, 93(3): 423–34.
- Bremmer, F. and Hoffmann, K.P. (1993) Pursuit related activity in macaque visual cortical areas MST and LIP is modulated by eye position. *Soc. Neurosci. Abstr.*, p. 1283.
- Brochic, P.R., Andersen, R.A., Snyder, L.H. and Goodman, S.J. (1995) Head position signals used by parietal neurons to encode locations of visual stimuli. *Nature*, 375(6528), 232–5.
- Burgess, N. (1995) A solvable connectionist model of immediate recall of ordered lists. In: Tesauro, G., Touretzky, D.S. and Leen, T.K. (Eds.), *Advances Neural Inform. Process. Syst.*, vol. 7. Cambridge, MA: MIT Press.
- Burnett-Stuart, G., Halligan, P.W. and Marshall, J.C. (1991) A newtonian model of perceptual distortion in visuo-spatial neglect. *NeuroReport*, 2: 255–257.
- Calvanio, R., Petrone, P.N. and Levine, D.N. (1987) Left visual spatial neglect is both environment-centered and body-centered. *Neurology*, 37: 1179–1181.
- Caramazza, A. and Hills, A.E. (1990) Spatial representation of words in the brain implied by studies of a unilateral neglect patient. *Nature*, 346: 267–269.
- Cohen, J.D., Farah, M.J., Romero, R.D. and Servan-Schreiber, D. (1994) Mechanisms of spatial attention: the relation of macrostructure to microstructure in parietal neglect. *J. Cogn. Neurosci.*, 6(4): 377–387.
- Colby, C.L. (1998) Action oriented spatial reference frames in cortex. *Neuron*, 20: 15–24.
- Colby, C.L. and Duhamel, J.R. (1993) Ventral intraparietal area of the macaque: anatomic location and visual response properties. *J. Neurophysiol.*, 69(3): 902–914.
- Deneve, S. and Pouget, A. (1998) Neural basis of object-centered representations. In: Jordan, M.I., Kearns, M.J. and Solla, S. (Eds.), *Advances in Neural Information Processing Systems*, vol. 10. MIT Press, Cambridge MA: MIT Press.
- Drain, M. and Reuter-Lorenz, P.A. (1997) Object-centered neglect for letters: do informational asymmetries play a role? *Neuropsychologia*, 35(4), 445–456.
- Driver, J., Baylis, G.C., Goodrich, S.J. and Rafal, R.D. (1994) Axis-based neglect of visual shapes. *Neuropsychologia*, 32(11): 1353–1365.
- Duhamel, J.R., Bremmer, F., BenHamed, S. and Graf, W. (1997) Spatial invariance of visual receptive fields in parietal cortex neurons. *Nature*, 389(6653): 845–848.
- Farah, M.J., Brunn, J.L., Wong, A.B., Wallace, M.A. and Carpenter, P.A. (1990) Frames of reference for allocating attention to space: evidence from the neglect syndrome. *Neuropsychologia*, 28(4), 335–347.
- Field, P.R. and Olson, C.R. (1994) Spatial analysis of somatosensory and visual stimuli by single neurons in macaque area 7B. *Soc. Neurosci. Abstr.*, vol. 20(1), p. 317.12.
- Galletti, C. and Battaglini, P.P. (1989) Gaze-dependent visual neurons in area V3a of monkey prestriate cortex. *J. Neuroscience*, 9: 1112–1125.
- Goldberg, M.E., Colby, C.L. and Duhamel, J.R. (1990) Cold Spring Harbor Symposium on Quantitative Biology. *Representation of visuomotor space in the parietal lobe of the monkey*, 55: 729–739.
- Goodale, M.A. and Milner, A.D. (1990) Separate visual pathways for perception and action. *Trends in Neuroscience*, 15: 20–25.

- Goodman, S.J. and Andersen, R.A. (1990) Algorithm programmed by a neural model for coordinate transformation. In: *International Joint Conference on Neural Networks*.
- Graziano, M.S., Yap, G.S. and Gross, C.G. (1994) Coding of visual space by premotor neurons. *Science*, 266, 1054–1057.
- Green, D.M. and Swets, J.A. (1966) Signal detection theory and psychophysics. John Wiley & Sons.
- Halligan, P.W. and Marshall, J.C. (1989) Line bisection in visuo-spatial neglect: disproof of a conjecture. *Cortex*, 25: 517–521.
- Heilman, K.M., Watson, R.T. and Valenstein, E. (1985) Neglect and related disorders. In: Heilman, K.M. and Valenstein, E. (Eds.), *Clinical Neuropsychology*. New York: Oxford University Press, pp. 243–294.
- Karnath, H.O., Schenkel, P. and Fischer, B. (1991) Trunk orientation as the determining factor of the 'contralateral' deficit in the neglect syndrome and as the physical anchor of the internal representation of body orientation in space. *Brain*, 114: 1997–2014.
- Karnath, H.O., Christ, K. and Hartje, W. (1993) Decrease of contralateral neglect by neck muscle vibration and spatial orientation of trunk midline. *Brain*, 116, 383–396.
- Kinsbourne, M. (1987) Mechanisms of unilateral neglect. Pages 69–86 of: Jeannerod, M. (Ed.), *Neurophysiological and Neuropsychological Aspects of Spatial Neglect*. North-Holland.
- Koch, C. and Ullman, S. (1985) Shifts in selective visual attention: towards the underlying neural circuitry. *Hum. Neurobiology*, 4(4), 219–27.
- Ladavas, E. (1987) Is the hemispatial deficit produced by right parietal lobe damage associated with retinal or gravitational coordinates? *Brain*, 110: 167–180.
- Ladavas, E., Pesce, M.D. and Provinciali, L. (1989) Unilateral attention deficits and hemispheric asymmetries in the control of visual attention. *Neuropsychologia*, 27(3): 353–366.
- Ladavas, E., Petronio, A. and Umiltà, C. (1990) The deployment of visual attention in the intact field of hemineglect patients. *Cortex*, 26: 307–317.
- Mazzoni, P. and Andersen, R.A. (1995) Gaze coding in the posterior parietal cortex. In: Arbib, M.A. (Ed.), *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press, pp. 423–426.
- Mishkin, M., Ungerleider, L.G. and Macko, K.A. (1983) Object vision and spatial vision: two cortical pathways. *Trends in Neurosci.*, Oct: 414–417.
- Mozer, M.C. and Behrmann, M. (1990) On the interaction of selective attention and lexical knowledge: a connectionist account of neglect dyslexia. *J. Cogn. Neurosci.*, 2(2): 96–123.
- Mozer, M.C., Halligan, P.W. and Marshall, J.C. (1997) The end of the line for a brain-damaged model of hemispatial neglect. *J. Cogn. Neurosci.*, 9(2): 171–190.
- Olson, C. and Gettner, S. (1998) Impairment of object-centered vision following lesion of macaque posterior parietal cortex. *Soc. Neurosci. Abstr.*, p. 449.10.
- Olson, C.R. and Gettner, S.N. (1995) Object-centered direction selectivity in the macaque supplementary eye. *Science*, 269: 985–988.
- Pizzamiglio, L., Frasca, R., Guariglia, C., Incoccia, C. and Antonucci, G. (1990) Effect of optokinetic stimulation in patients with visual neglect. *Cortex*, 26: 535–540.
- Poggio, T. (1990) A theory of how the brain might work. *Cold Spring Harbor Symposium on Quantitative Biology*, 55: 899–910.
- Poggio, T. and Girosi, F. (1990) Regularization algorithms for learning that are equivalent to multilayer networks. *Science*, 247: 978–982.
- Pouget, A. and Driver, J. (1999) Visual neglect. In: Wilson, R. and Keil, F. (Eds.), *MIT Encyclopedia and Cognitive Sciences*. Cambridge: MIT Press.
- Pouget, A. and Sejnowski, T.J. (1995) Spatial representations in the parietal cortex may use basis functions. In: Tesauro, G., Touretzky, D.S. and Leen, T.K. (Eds.), *Advances in Neural Information Processing Systems*, vol. 7. Cambridge, MA: MIT Press.
- Pouget, A. and Sejnowski, T.J. (1997a) A new view of hemineglect based on the response properties of parietal neurons. *Philosoph. Transac. Roy. Soc. Series B*, 352: 1449–1459.
- Pouget, A. and Sejnowski, T.J. (1997b) Spatial transformations in the parietal cortex using basis functions. *J. Cogn. Neurosci.*, 9(2): 222–237.
- Pouget, A. and Sejnowski, T.J. (1999) Lesioning a basis function model of spatial representations in the parietal cortex: comparison with hemineglect. *Psycholog. Rev.*, submitted.
- Rubens, A.B. (1985) Caloric stimulation and unilateral visual neglect. *Neurology*, 35: 1019–1024.
- Salinas, E. and Abbott, L.F. (1995) Transfer of coded information from sensory to motor networks. *J. Neurosci.*, 15(10): 6461–74.
- Salinas, E. and Abbott, L.F. (1996) A model of multiplicative neural responses in parietal cortex. *Proc. Nat. Acad. Sci., USA*, 93: 11956–11961.
- Schlag-Rey, M. and Schlag, J. (1984) Visuomotor functions of central thalamus in monkey. I. Unit activity related to spontaneous eye movements. *J. Neurophysiol.*, 51(6), 1149–74.
- Sparks, D.L. (1991) Sensori-motor integration in the primate superior colliculus. *Seminars in the Neurosciences*, 3: 39–50.
- Stein, J.F. (1992) The representation of egocentric space in the posterior parietal cortex. *Behav. Brain Sci.*, 15(4): 691–700.
- Tipper, S.P. and Behrmann, M. (1996) Object-centered not scene-based visual neglect. *J. Experimental Psychology, Human Perception and Performance*, 22(5): 1261–1278.
- Treisman, A. and Gelade, G. (1980) A feature integration theory of attention. *Cogn. Psychol.*, 12: 97–136.
- Trotter, Y., Celebrini, S., Stricanne, B., Thorpe, S. and Imbert, M. (1992) Modulation of neural stereoscopic processing in primate area V1 by the viewing distance. *Science*, 257: 1279–81.

Van Opstal, A.J., Hepp, K., Suzuki, Y. and Henn, V. (1995) Influence of eye position on activity in monkey superior colliculus. *J. Neurophysiol.*, 74(4): 1593–1610.

Zipser, D. and Andersen, R.A. (1988) A back-propagation programmed network that stimulates responses properties of a subset of posterior parietal neurons. *Nature*, 331: 679–684.