



PERGAMON

Neural Networks 11 (1998) 1229–1240

Neural
Networks

1998 Special Issue

The maintenance of spatial accuracy by the perisaccadic remapping of visual receptive fields

Christian Quaia¹, Lance M. Optican, Michael E. Goldberg^{2*}

Laboratory of Sensorimotor Research, National Eye Institute Bethesda, MD 20892, USA

Received and accepted 6 May 1998

Abstract

Humans and monkeys can direct their eyes to the spatial location of briefly flashed targets even when a saccade intervenes between the stimulus flash and the saccade to acquire its location. It had been proposed that the oculomotor system performs this task by resorting to a supramaxillary representation of space. In this paper we review neurophysiological and clinical data suggesting that the brain can use a different strategy that does not require an explicit supramaxillary representation of targets. We propose and implement a simple neural model that can keep track continuously of the location of saccade targets in eye-centered coordinates. Finally, based on recent data, we argue that such a neural mechanism is in fact used to keep track not only of saccade targets but of the location of salient areas of the visual scene in general. Published by Elsevier Science Ltd.

Keywords: Model; Saccades; Saliency; Lateral intraparietal cortex (LIP); Frontal eye fields (FEF); Superior colliculus (SC)

1. Introduction

Saccades are rapid eye movements made to a visual target. The brain chooses a target in the visual periphery and drives the eyes as rapidly as possible so that at the end of the movement the target that was in the peripheral visual field is now on the fovea. Early models of the saccadic system assumed that it used the retinal location of the target to calculate the desired saccade (for a review see Westheimer, 1989). However, Hallett and Lightstone, using the double-step task (Fig. 1), showed that humans can make accurate saccades to targets that briefly appeared before an intervening saccade (Hallett and Lightstone, 1976). Similarly, Mays and Sparks (1980a), Mays and Sparks (1980b) reported that monkeys can make accurate eye movements to a flashed target even when the eyes are artificially displaced (using electrical stimulation) after the target has disappeared. Thus, it can be concluded that the brain can accurately keep track of the location of areas of interest (like targets for a subsequent saccade).

To account for this ability to make accurate saccades in the double-step task, Robinson (1975) proposed that

saccadic targets are stored in spatial coordinates, as opposed to retinal coordinates (Fig. 2). Robinson noted that, if it is assumed that the head is stationary, the position of the target in space can be replaced by the position of the target relative to the head (TreH), which can be computed simply by summing the position of the eye in the head at the time the target appeared (EreH) to the retinal location of the target (TreE). To compute the displacement necessary to foveate the target it is then sufficient to compute the difference between TreH and the current position of the eye in the head (EreH). Zeroing this difference is then sufficient to guarantee that the target is accurately foveated, regardless of intervening movements of the eyes. The simplicity of this solution (even though other signals had to be introduced to account for movements of the head or of the body in general) made it very appealing and induced several investigators to look for the neural correlates of the various signals postulated in the Robinson model. Thanks to those efforts, a large body of data is now available about the pattern of activity present in several brain areas [e.g. superior colliculus (SC), frontal eye fields (FEF) and lateral intraparietal cortex (LIP)] during double-step saccades. However, the explicit representation of target position in space (or even relative to the head) has not been found.

In this paper we will argue that there is no need for such an explicit representation and we will review the evidence about the existence of a neural mechanism that can keep

* Corresponding author. Tel: 301-496-1060; Fax: 301-402-0511; e-mail: meg@lsc.nei.nih.gov

¹ Also at DEEI, Università degli Studi di Trieste, Trieste, Italy.

² Also at the Department of Neurology, Georgetown University School of Medicine, Washington, DC 20007, USA.




Fig. 1. The double-step task. Subject fixates A. A disappears and simultaneously B appears. B remains on for 100 ms, then disappears and C flashes on for 50 ms. The subject has to make first a saccade to the location B (saccadic vector \overline{AB}), and then to the location C (\overline{BC}). If the saccadic system were to use only retinotopic information, the second saccade would direct to C' (saccadic vector $\overline{BC'} = \overline{AC}$) and not to the desired location C.

track continuously of the location of saccade targets without resorting to a suparetinal representation of space. Based on recent data (Gottlieb et al., 1998), we will argue that such a neural mechanism is in fact used to update not only sensorimotor maps (such as a map of saccade targets) but saliency maps (i.e. maps that store the location of meaningful areas of the visual scene) in general. We will also present a model that describes a possible implementation of such a mechanism.

2. Neuronal activity in the double-step task

Mays and Sparks (1980a) were the first to investigate the activity of single neurons in the double-step task. Recording from the intermediate layers of the superior colliculus, they discovered that neurons with visual and presaccadic responses appropriate for the second saccade begin to discharge immediately after the first saccade of a double step, even though the visual target never entered their receptive

field. They called cells that responded under these circumstances 'quasi-visual cells'.

Although the quasi-visual cells showed that the superior colliculus had access to extraretinal signals, it was not clear at what level information about eye movements or eye position entered the oculomotor system. In fact, quasi-visual cells, located in the intermediate layers of the colliculus, had both movement and visual activity, and it was thus impossible to distinguish between an interaction of eye movements and visual processing, or eye movements and motor output processing. Fortunately this distinction could be made in the frontal eye fields, where the neurons that discharge before visually guided saccades can be divided into three different classes (Bruce and Goldberg, 1985). Visual neurons respond to the onset of visual stimuli whether or not they are saccade targets, but do not discharge before memory guided saccades or learned saccades in total darkness. Movement neurons have little or no response to the appearance of a visual stimulus, even one that is a saccade target, but discharge before all purposive saccades of a certain amplitude and direction, whether or not those saccades are made to a visual target. Visuomovement neurons have both visual and saccadic responses, and tend to discharge most before visually guided saccades.

All three types of neurons discharge before the second saccade of a double-step saccade task if they would ordinarily discharge before that saccade in a simple visually guided saccade task (Goldberg and Bruce, 1990). It is remarkable that visual neurons without any movement activity respond in this situation, because the stimulus never appeared in the neuron's receptive field (as determined in a fixation task) (Fig. 3). Control experiments showed that those neurons do not discharge when the stimulus is flashed but the monkey does not make the first saccade, nor do they discharge if the monkey makes the first saccade but the second target was not flashed. This demonstrates that even neurons that are not tightly linked with the execution of a saccade (they normally discharge

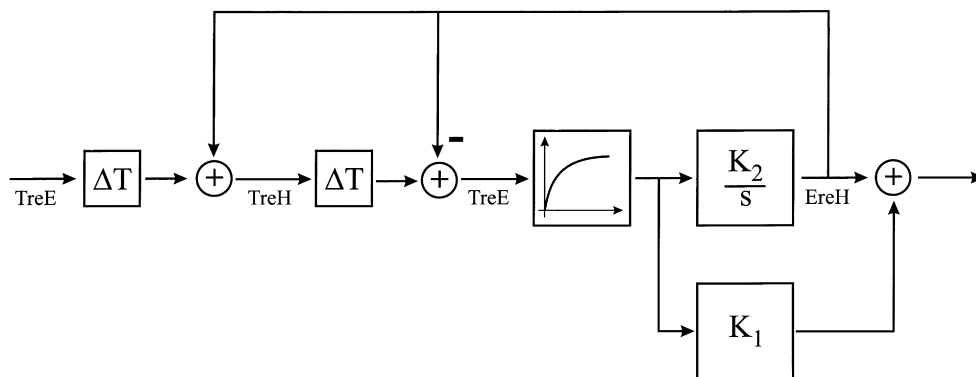


Fig. 2. Simplified description of the Robinson model (Robinson, 1975). The visual input enters the brain through the retina, and the location of a target is thus encoded relatively to the eye (TreE). After a time ΔT (due to transmission delays and possibly visual processing) the visual input reaches an unspecified brain area where the location of the target relative to the head (TreH) is computed by summing the location of the eye relative to the head (EreH) to TreE. The targets are memorized in this map, and to foveate any given target it is then sufficient to bring to zero, using a negative feedback loop, the difference between TreH and the current position of the eyes (EreH). The velocity signal to drive the eyes is generated applying a non-linear filter to the motor error (TreE) signal; this signal is then integrated, and appropriately weighted, to obtain an estimate of the position of the eyes in the head (EreH).

whether or not a saccade is made) can specify the location of a target of interest despite an intervening saccade. In other words, these neurons are able to keep track of the location of regions of interest in the visual scene, regardless of intervening saccades.

From these results it can be concluded that visual neurons in the frontal eye field describe the location of a relevant stimulus in an oculocentric coordinate system. After a movement the origin of the coordinate system, which initially coincided with the spatial location of the fovea at the

time the stimulus appeared, coincides with the spatial location of the current center of gaze. Nonetheless, the information carried by the neuron is unchanged: these neurons signal the presence of a possible saccade target a certain distance and direction from the fovea. Ordinarily this output vector corresponds to the vector from the fovea to the image of the target on the retina. However, an intervening saccade, such as the first saccade in the double-step task, changes the oculomotor response to the target (i.e. the eye movement needed to point the fovea toward the location of the target).

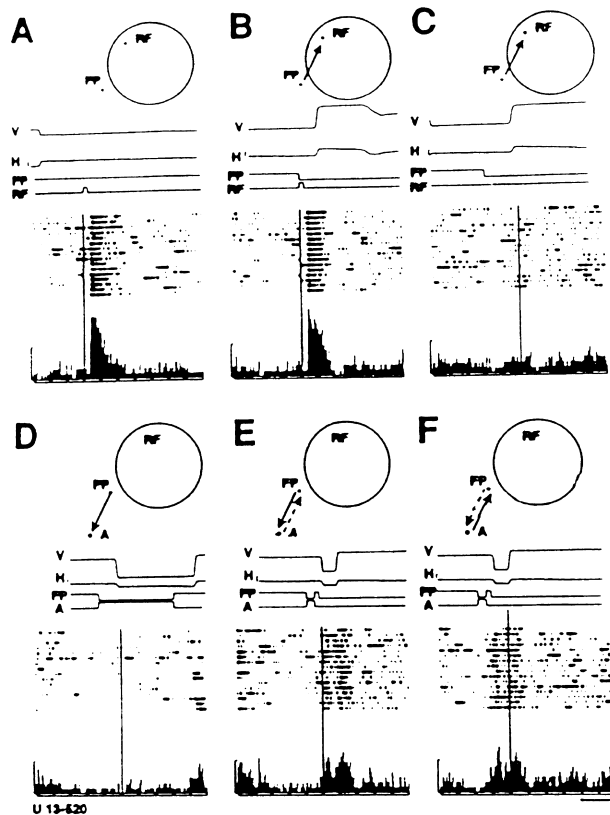


Fig. 3. Response of a frontal eye field visual neuron in the double-step task. Each panel has a cartoon showing the location of the receptive field at the time of fixation (RF), the fixation point (FP), and the saccades. Beneath the cartoon are sample horizontal (H) and vertical (V) eye position traces, and a stimulus trace showing time of stimulus appearance. Neuronal responses are shown beneath as rasters and histograms. In the raster each dot is a cell discharge and each line a single trial. Successive lines are synchronized on a single event which occurred at the time of the vertical line. The histogram sums the raster above. (A) Response to the appearance of the stimulus in the RF when the monkey does not make a saccade to the stimulus. (B) Response when the monkey makes a saccade to the RF stimulus. (C) No response when the monkey makes the saccade when the stimulus has not recently appeared. This establishes that the neuron is a visual cell without any movement activity. (D) No response when the monkey makes a saccade away from the RF to A. There is no response either to target or saccade. (E) and (F) Double-step task. The saccade target appears at A, then the FP flashes again briefly, and disappears. The monkey makes two subsequent saccades, to A and then back to RF. A is situated so that the FP is in the receptive field when the monkey looks at A. (E) Activity synchronized on the first saccade. (F) Activity synchronized on second saccade. The cell discharges at the time of the first saccade until the second. Reproduced with permission from Goldberg and Bruce (1990).

3. Pre- and perisaccadic remapping

The presence of such a coordinate shift mechanism provides a neural substrate for the ability of subjects to perform double-step paradigms. However, these results cannot delineate the mechanism that makes such change of coordinates possible. In fact, at least two solutions are possible. First, the location of relevant targets could be stored in spatial coordinates in some other brain area, and the activity recorded in the frontal eye field and in the superior colliculus could simply reflect the difference between this signal and the current location of the eyes, in agreement with what Robinson (1975) proposed (see Fig. 2).

Alternatively, Goldberg and Bruce (1990) proposed that the saccade needed to foveate the spatial location of the target could be computed by subtracting from the vector describing the original retinal location of the target the vector of the intervening saccade (Fig. 4). So, by dynamically calculating the postsaccadic vector, the brain could accurately localize the saccade target in space using a continuously updated gaze-centered vector, without ever requiring an explicit representation of target position in space. The saccade generator could then be organized as proposed by Jürgens et al. (1981), comparing the desired displacement signal (i.e. TreE) with the displacement of the eyes since the beginning of the saccade, which can be obtained by integrating the eye velocity signal over the course of the current saccade (Fig. 5; note that the integrator must be reset before the beginning of each saccade).




Fig. 4. Vector subtraction hypothesis. In the double-step task (Fig. 1), after the first saccade from A to B the subject must foveate C. The saccadic vector describing the displacement from B to C (\overline{BC}) can be computed by simply subtracting from the vector, which described the retinotopic location of the target at the time it was flashed, the vector describing the intervening saccade, i.e. \overline{AB} .

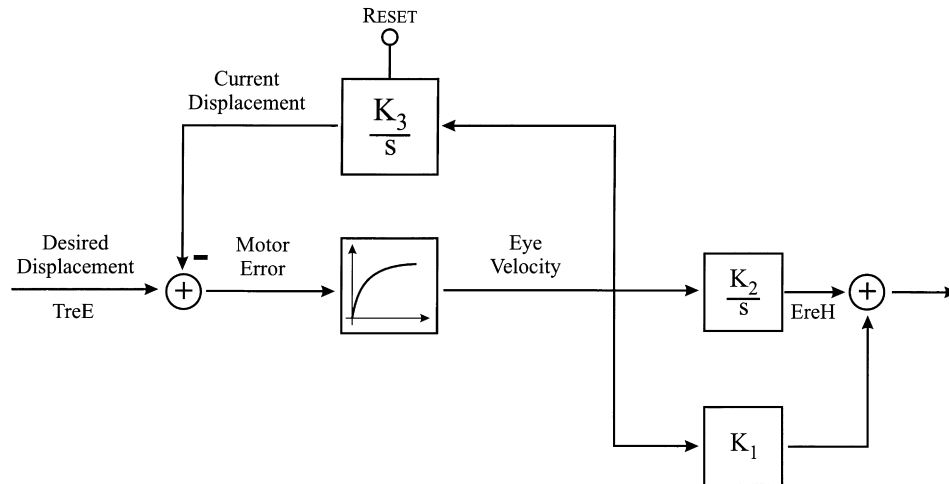


Fig. 5. Schematic description of the model of Jürgens et al. (1981). In this case the desired displacement of the eyes is computed for each movement, and then a negative feedback loop is used to displace the eyes by that same amount. This is accomplished generating an eye velocity signal, which is then integrated to obtain an estimate of the displacement of the eyes since the beginning of the saccade (thus the integrator must be reset at the beginning of each saccade). The difference between the desired displacement and the current displacement is then the signal that must be brought to zero. In this model the position of the eyes relative to the head (EreH) is still computed, but it is used only to keep the eyes still at the end of the movement, and not to determine the amplitude of the saccade.

The double-step task is a laboratory curiosity, but it illustrates a more general problem. The latency of frontal and parietal visual areas is at least 60 ms (Bushnell et al., 1981; Goldberg and Bushnell, 1981). This means that every time we move our eyes the information in the brain is inaccurate for the purposes of localization for at least 60 ms. If we assume that humans ordinarily make three saccades a second during visual exploration, then the cortical representation of the visual world would be inaccurate for

the purpose of spatial localization at least 180 ms out of every second, or roughly 20% of the time! Such an inaccuracy in spatial localization, and by inference for spatial behavior, would have very adverse consequences for evolution.

The coordinate shift mechanism first discovered in double-step experiments provides a solution to this behavioral quandary. Every time a monkey makes a saccade, there is a shift of the visual world on the retina into the

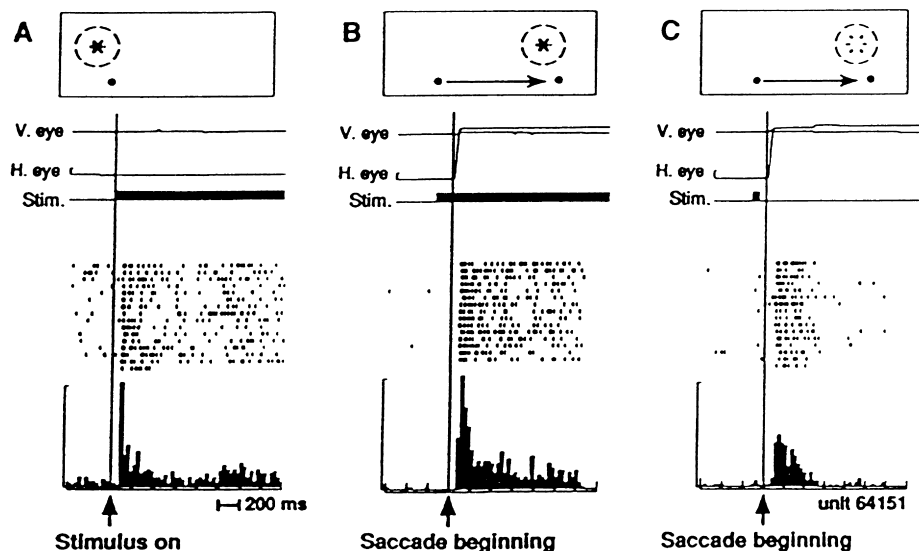


Fig. 6. Remapping visual memory trace in a neuron in LIP. Each panel contains a cartoon showing the spatial locations of the receptive field (circle with stimulus), fixation point (small filled circle) and saccade (arrow). Eye and stimulus position traces and neuronal responses as in Fig. 3. (A) Response to target in visual receptive field. There is an on-response and a smaller tonic visual response. Activity synchronized on stimulus appearance. (B) Response when monkey makes a saccade bringing stimulus into receptive field. The cell discharges after the stimulus is brought into the receptive field of the saccade. Note there is no response when the stimulus appears, showing that the neuron does not respond to the stimulus when it appeared outside the receptive field. Activity synchronized on beginning of saccade. (C) Activity when stimulus flashes outside receptive field before saccade, but saccade brings spatial location of stimulus into receptive field. Activity synchronized on saccade beginning. The neuron did not respond when the monkey made the saccade without the stimulus being present (not shown). Reproduced with permission from Duhamel et al. (1992a).

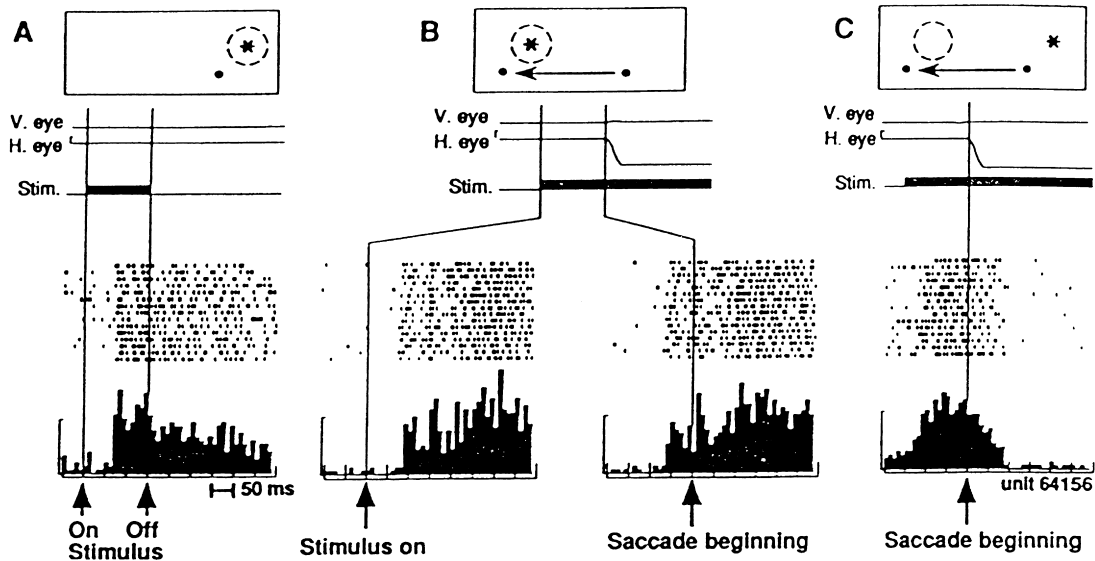


Fig. 7. Predictive visual response in LIP neuron. Cartoons and diagrams as in Fig. 6. (A) Response to the appearance of stimulus in receptive field. Note that although the stimulus appears for 600 ms, the response remains much longer. Activity synchronized on stimulus appearance. (B) Response when monkey makes saccade bringing stimulus into receptive field. Activity synchronized on stimulus appearance (left) and saccade beginning (right). Note that response occurs before saccade, although the latency from the stimulus appearance is longer than when the stimulus appears in the receptive field. (C) Truncation of response when monkey makes saccade bringing stimulus out of receptive field. Reproduced with permission from Duhamel et al. (1992a).

coordinates of the new fixation. The cell pictured in Fig. 6 has a retinotopic receptive field. It responds when a stimulus is in the receptive field but not to the same stimulus in the same spatial location when the monkey is fixating a different place and has therefore moved the retinal receptive field away from the stimulus. However, under one circumstance it will respond to a stimulus that is not in its retinotopic receptive field. When the monkey makes a saccade that brings the spatial location of the stimulus into the retinotopic receptive field of the neuron, the cell will respond to the stimulus even though it flashed only before the saccade, never actually entered the retinal receptive field as determined in a fixation task (Fig. 6C), and no second saccade to its location was ever made. The response is not merely a postsaccadic response to the saccade: there is no postsaccadic activity if the monkey makes a similar saccade that does not move the location of a relevant stimulus into the receptive field.

This saccade-related effect on the response of a visual neuron has an inverse. There is a class of LIP neurons with tonic responses to transiently flashed targets. Thus, the disappearance of the target is not enough to turn off the cells' response (Fig. 7A), which is reasonable because the spatial location of the stimulus is still described by the same vector. However, when the monkey makes a saccade that removes the stimulus from the receptive field the response is truncated (Fig. 7C). We propose that this happens because the location of the stimulus relative to the eye no longer coincides with that neuron's output vector, and if the cell were to continue to discharge it would provide a spatially inaccurate signal.

This shift of receptive field can occur even before the saccade occurs. In Fig. 7B one such neuron is shown; note that the neuron consistently (see rasters) discharges before the

saccade begins. However this neuron does not respond to the stimulus in that retinal location when the monkey does not execute a saccade that brings the location of the stimulus into the receptive field (as determined in a fixation task). Because the neurons respond as if they predict that the location of the stimulus will enter their receptive fields, these responses have been called predictive visual responses.

It is important to emphasize that the predictive response is entirely unrelated to any saccade that the monkey might make to the stimulus whose response is remapped. Although the original observations (Duhamel et al., 1992b) were made




Fig. 8. Adjusted saccade latencies plotted in order for 80 visual and visuomovement neurons in FEF. The adjusted latency is the visual latency minus the latency from the saccade when the monkey brings the stimulus into the receptive field. Any neuron whose adjusted latency is less than zero has a predictive visual response. The horizontal lines show 0 ms and the time at which the longest saccade begins (-60 ms). Reproduced with permission from Umeno and Goldberg (1997).

in monkeys that had previous experience with the double step task, the finding has been replicated in LIP (Gottlieb et al., 1998) and SC (Walker et al., 1995) in monkeys that had never performed a double-step task, and in FEF in monkeys who were preparing a second saccade away from the location of the second target (Umeno & Goldberg, 1997).

The time difference (latency) between the saccade that brings the location of the stimulus into the receptive field of the neuron and the onset of the predictive response spans a wide range (Fig. 8 shows the distribution of latencies in FEF). When the monkey makes a saccade that brings a stable stimulus into its receptive field, one would expect the cell to respond with a visual latency. Neurons that discharge even though the target was only briefly flashed must then rely on some neural mechanism that does not require postsaccadic retinal information. The presence of neurons that discharge with a latency shorter than the visual latency means that this neural mechanism is faster than the visual afferents; more interestingly, the presence of cells that discharge even before the saccade excludes eye position signals as being an indispensable component of this mechanism and makes a Robinson-like scheme very difficult to implement.

4. Clinical deficits in the double-step task

If the brain can perform the double-step task because it can compensate for the first saccade, then one would expect

that parietal damage would result not in a spatial deficit per se, but rather in a deficit in compensation for the saccade. This phenomenon is observed in humans with parietal damage (Duhamel et al., 1992b; Heide et al., 1995). A patient with a large right frontoparietal lesion was studied in single- and double-step saccade tasks (Duhamel et al., 1992b). The lesion had occurred 14 years before she was studied. Her leftward saccades were hypometric and had longer latencies than her rightward ones, and this difference was independent of where the eye was in the orbit at the beginning of the saccade. The deficit was not a simple spatial one: a stimulus at a given spatial location would be acquired inaccurately when it was approached by a leftward saccade, and accurately when it was approached by a rightward one.

Despite her deficit, the patient could perform the double-step task well when she made the first saccade in a rightward direction, and then had to make a second saccade in the leftward direction, to a target that only appeared in the right visual field (Fig. 9). She could clearly compensate for the rightward saccade, and her subsequent leftward saccade was only slightly less accurate than her visually guided leftward saccades. When the first saccade was in the leftward direction the saccade was, as expected, less accurate and of longer latency than saccades to the right. However, she was then unable to make the second saccade, a saccade in the unaffected direction, to a stimulus that had appeared in her unaffected field. Her deficit was her inability

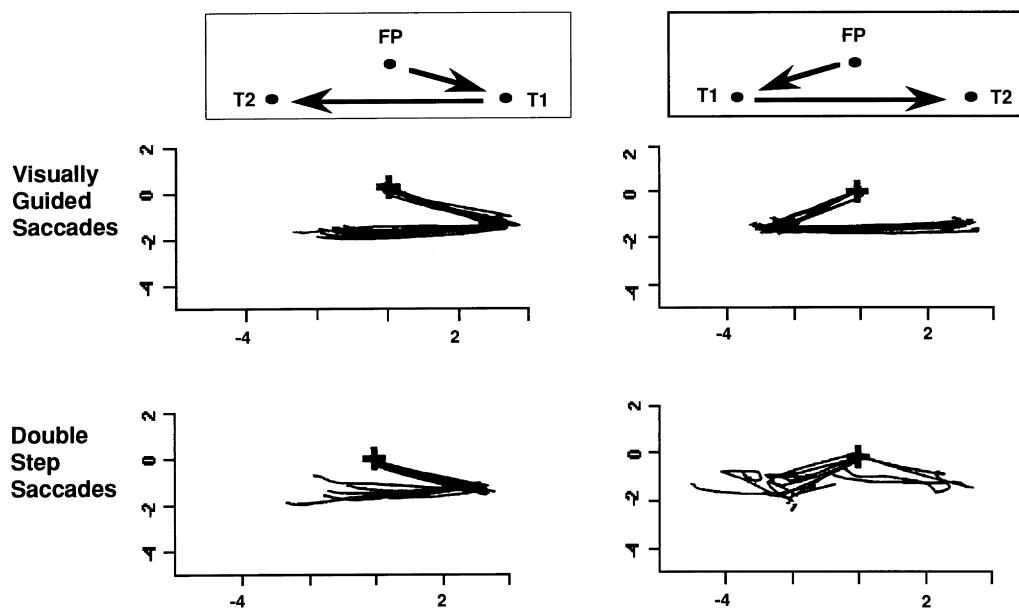


Fig. 9. Performance of patient with frontoparietal lesion in visually guided and double-step saccade tasks. Each panel shows eye movements plotted in the X–Y plane. Top row: the cartoon above shows the arrangement of the fixation point (FP) and two saccade targets (T1 and T2), and the saccades made to them. In the left column the first saccade is made into the normal field, and the second into the affected field. Middle row: each stimulus remained on for 500 ms, and the subject made two successive visually guided saccades. The leftward saccades are all slightly inaccurate. Bottom row: in each panel T1 appears for 100 ms and T2 for 80 ms. In the right figure the saccades into the normal field are accurate, and the saccades into the affected field slightly more inaccurate than in the visually-guided case. Right panel: when T1 appears in the affected field, the subject usually makes a slightly inaccurate first saccade, although sometimes she makes a saccade directly to the target in the normal field. When she makes the saccade into the affected field, she cannot make the second saccade into the normal field, despite the fact that the stimulus is in the normal field and the desired saccade in the normal direction. Reproduced with permission from Duhamel et al. (1992b).

to compensate for the saccade into the affected field. This is not a memory deficit or a spatial deficit. The patient occasionally made saccades directly to the second target location, indicating that she both saw and remembered its location. Her problem was that she could not calculate the change in target position relative to eye position that was caused by the first saccade. The likely explanation for this was that she lacked the mechanism that remapped the visual signal into the coordinates of the fixation that followed the first, leftward saccade. Although this was only a single case report, the phenomenon has been replicated in a large number of patients and was exhibited only by patients whose lesion included the right posterior parietal cortex and not patients with damage limited to frontal cortex (Heide et al., 1995).

5. A computational model of the remapping process

The physiological and clinical results described above make it clear that visual processing in the oculomotor system provides the saccade generator with a signal that allows an accurate foveation of the target, irrespective of intervening eye movements. Many workers have postulated that such a signal arises through a mechanism that involves the computation of absolute target position in suparetinal coordinates (Mays and Sparks, 1980b; Schlag and Schlag-Rey, 1991; Zipser and Andersen, 1988), and then the subtraction between this signal and the position of the eyes. However, because of the predictive nature of the responses, this position signal must have predictive capabilities, and cannot just be the efference copy of the current eye position signal. Up to now, neither the representation of the target in space nor a predictive eye position signal have been reported. In contrast, we show here how the remapping observed in LIP, FEF and SC can be reproduced by a network that uses only signals that have already been observed during neurophysiological recordings.

Starting from the hypothesis that the remapping of activity is performed in LIP whereas the remapping observed in FEF and SC (which receive projections from LIP) is caused by what happens in LIP, we have modeled six classes of cells, three in LIP and three in FEF. To simplify the presentation we will illustrate the connectivity of our model in three separate steps. First we will show how the model predicts the patterns of activity observed in response to a flashed target; second, we will elucidate how the predictive remapping arises; and third, we will show how the wide range of latencies can be accounted for.

5.1. Structure of the model

In Fig. 10 we present the basic connectivity of the model that we used to replicate the visual responses of LIP and FEF neurons. We hypothesize that the LIP phasic cells (LIPP), i.e. the cells that produce a burst of activity when the target appears but have no sustained response, receive

the visual input from higher visual areas. We propose that the visual input is supplied directly to the phasic cell and to an interneuron, which acts as a low pass filter. The output of this interneuron then is subtracted off, at the dendritic level, from the visual input. This circuit forms a high pass filter, and guarantees that when a target appears only a transient response is observed, with the activity growing very rapidly and then decaying more slowly. We propose that the interneuron acts as a dendritic shunt, because otherwise the extinction of the target would cause a transient strong inhibition of the cell (thus making it difficult to perform a remapping to a cell that was already active; however, it is actually not known what happens under these conditions). Nonetheless, appropriate weights and time constants could probably be chosen to make this dendritic operation unnecessary.

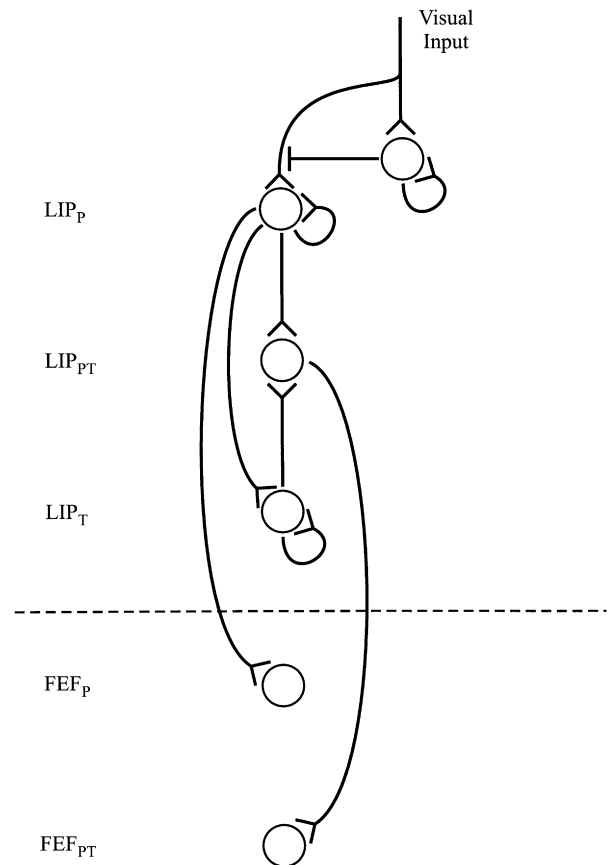


Fig. 10. Model connectivity sufficient to replicate visual responses. The visual input (which is assumed to vary in a stepwise manner) is processed by an interneuron that behaves as a low pass filter ($T_c = 5$ ms). The output of the interneuron is subtracted from the visual input, and this signal (which is a high pass of the visual input) is supplied to the LIP phasic cells (LIP_P). The output of the LIP_P cell is a low pass ($T_c = 90$ ms) of its input. The output of the LIP_P cell is then supplied to a FEF phasic cell (FEF_P), to a LIP phasic-tonic cell (LIP_{PT}), and to a LIP tonic cell (LIP_T). The LIP_T cell behaves as a low pass filter with a high time constant ($T_c = 3$ s), and thus essentially integrates its input. The LIP_T cell projects to the LIP_{PT} cell, which in turn drives the FEF phasic-tonic cell (FEF_{PT}). The LIP_{PT}, FEF_P and FEF_{PT} cells are all modeled as low pass filters of their inputs ($T_c = 5$ ms).

The LIP_P cell then acts as a low pass on its input, and passes it on to the phasic cells in the FEF (FEF_P). We did not supply any direct visual input to the FEF, but this could be introduced without affecting in any way the behavior of the model. The LIP_P signal is also passed on to the LIP tonic cells (LIP_T), which are low pass filters with a fairly long time constant, i.e. they essentially act as an integrator. Thus, these cells have a sustained firing rate if a target appeared in their receptive field but do not burst at its onset. Both the LIP_P and the LIP_T cells then project to the LIP phasic-tonic cells (LIP_{PT}), which are the most commonly reported cells in LIP. The phasic activity in these cells is thus determined by the input they receive from the phasic cells, whereas the sustained activity reflects the activity in the tonic cells. This arrangement implies that if cells receive smaller visual input (for example because the stimulus is not in the center of the receptive field) both the phasic and the tonic components will be smaller. This characteristic is supported by experimental evidence (Barash et al., 1991b). It will become clear later on that all these cells are necessary to perform the remapping in LIP, and the LIP_{PT} cells alone can then induce the remapping observed in the SC and in the FEF. To calibrate the strength of the various connections we used quantitative data from recordings in LIP_{PT} cells projecting to the intermediate layers of the SC (Paré and Wurtz, 1997).

Suppose that a target is initially flashed 30° to the right of the fovea, and that the subjects then makes a 10° rightward saccade. Based upon the experimental evidence we have reviewed above, when the target flashes the LIP cells with

a receptive field 30° to the right of the fovea should discharge; then, as early as 80 ms before the onset of the saccade, the cells whose receptive field is centered 20° to the right of the fovea should become active, and by the end of the saccade the first population of cells should become inactive (so that an unambiguous representation of the target location is available).

In Fig. 11 we illustrate the basic connectivity that enables our model to produce such predictive remapping. We propose that the signal that carries the information about the impending saccade is represented by the output of the FEF movement cells (FEF_{MOV}), that start discharging approximately 100 ms before saccade onset (Bruce and Goldberg, 1985). Consider now again our specific example: the 30° LIP_{PT} cells are active, and 100 ms before saccade onset the 10° FEF_{MOV} cells start discharging. Every time that these two populations of cells are active at the same time, the LIP 20° cells must be activated; it is thus sufficient to multiply, at the level of the dendrites of the LIP 20° cells, the 30° LIP_{PT} signal by the 10° FEF_{MOV} signal. More generally, all the pairs (LIP_{PT} , FEF_{MOV}) whose receptive/movement fields difference is equal to 20° must be connected to the 20° LIP_P cells. We will explain later on how such a very specific connectivity could emerge trivially. Once the multiplication (or more in general any AND operation, i.e. the output has to be vigorous if and only if both inputs are active) is performed, the 20° LIP_P cell becomes active, and the remapping takes place. Because of the time course of the FEF_{MOV} signal the prediction is




Fig. 11. Model connectivity needed to produce the remapping. If a target has been presented 30° to the right and a 10° rightward saccade is executed, the 20° site will have to be activated. This is accomplished by multiplying the output of the 30° LIP_{PT} cell by the output of the 10° FEF_{MOV} cells, which start discharging 100 ms before the saccade begins. The multiplication operator could be replaced by any operator that guarantees that the output is strong if and only if both inputs are strong. To remove the activity from the 20° cells, a damping signal is provided, during the saccade, to all LIP_T cells.

in the same range observed by Duhamel et al. (1992a). Note that the remapping input must be supplied to the LIP_p cells, to guarantee that this signal is then integrated by the LIP_T cells.

As we already pointed out, and as confirmed by neural recording (see above), by the end of the first saccade the 30° LIP_T cells must be silenced. We propose that this is accomplished by sending, during each saccade, a non-specific damping signal to all the tonic cells in LIP. This signal will be sufficient to silence the 30° LIP_T cells, and it will not allow the 20° LIP_T cells to charge during the saccade, but, because the remapping input is activated before saccade onset and the LIP_p cells act as a low pass filter, after the end of the movement the 20° LIP_p cells will still be active, thus charging the 20° LIP_T cells and stabilizing the remapping.

The structure we have described so far can simulate the predictive remapping, but it does not account for the wide range of delays with which the remapping occurs in the various cells. To account for this characteristic, which has been reported in LIP, FEF and SC, we propose that only a very small subset of LIP_p cells receive the remapping

inputs, and that these cells are those that exhibit the earliest remapping. We proposed earlier that each LIP_p projects to a FEF_p cell (Fig. 10); accordingly, the FEF cells that have the shorter latencies are those that receive their input from the LIP cells with the shortest latencies. We propose that in turn, the FEF_p cells project to other LIP_p cells (i.e. not to the LIP_p cell that gave input to the FEF_p cell itself) having the same receptive field, thus creating a sort of ‘bucket brigade’ that spreads the activity throughout the LIP and FEF populations encoding the new target location (Fig. 12). Delays in the connections of a few milliseconds and the delay introduced by the low pass filters present in the system can easily reproduce the large span of latencies experimentally observed. We propose then that the LIP_{PT} cells that induce the remapping when a movement is about to be executed (i.e. those whose activity is multiplied by the FEF_{MOV} activity) are those with relatively long latencies. This arrangement guarantees that the system remaps the activity only once for each saccade, and it might be the reason for having such a wide range of latencies.




Fig. 12. Model connectivity to reproduce the different remapping latencies experimentally observed. We propose that each FEF_p cell projects to a LIP_p cell which is different from the one that drives the FEF_p cell itself. Because of the transmission delays and of the low pass filtering of the LIP_p cells, cells with different latencies arise. To avoid a crescendo of activity in the chain in response to a visual stimulus, the FEF_p input must be shunted by the visual input. However, to make this input effective for remapping the visual input must be shunted by the motor command.

There are two major problems that need to be solved for this scheme to work. First, the LIP cells (as well as the FEF cells), all have the same latency when a visual target is presented in their receptive field; thus, they must all receive a direct visual input. However, this would result in a crescendo of activation because the first cell in the chain would receive only the visual input, but the second cell would receive its visual input plus the one relayed by the FEF_P cell in the first chain, and so on. A solution to stabilize the gain is to have the visual input to the LIP_P neurons shunt the input from the FEF_P cells. The second problem follows from the solution of the first: if the projection from the FEF is gated by the visual input, how can the remapping occur if the visual input is not removed? The solution to this problem is actually very simple: the visual input must be shunted by the motor command (i.e. by the FEF_{MOV} neurons). Note that this solution will make the system less responsive to targets flashed from around 100 ms before saccade beginning to saccade end. This presaccadic decrease in sensitivity has been demonstrated for neurons in LIP (Kusunoki et al., 1994).

5.2. Simulations

We now report the pattern of activity observed in the various units of the network described above during a

simulation of the model. More extensive simulations will be presented in another paper. We decided to simulate the following experimental conditions: at time 50 ms a target is briefly flashed 30° to the right of the fovea, and at time 1 s a 10° rightward saccade is performed. In Fig. 13A we show the pattern of activity observed in the 30° LIP_{PT} cell: the cell bursts approximately 60 ms after the onset of the target, and then its activity decays but does not vanish. This replicates both qualitatively and quantitatively, recordings in LIP_{PT} neurons (e.g. see Barash et al., 1991a; Duhamel et al., 1992a; Paré and Wurtz, 1997). When the saccade is produced, the location of the (now extinguished) target is removed from the receptive field of the neuron, and its activity goes quickly to zero (Fig. 13A). As previously pointed out, this is exactly what happens in LIP neurons (Fig. 7C).

We now analyze what happens to the 20° LIP_{PT} cells. We plotted the behavior of two cells, the one that exhibits the earliest remapping (Fig. 13B) and one that show remapping only much later (Fig. 13C). It is easy to see that a wide range of latencies are spanned. Finally we show what happens to the earliest 20° LIP_{PT} cell if the saccade is made to a stationary target as opposed to a flashed target. In this case the cell gets a reafferent visual response, and its discharge is increased (compare part B with part D). Note the striking

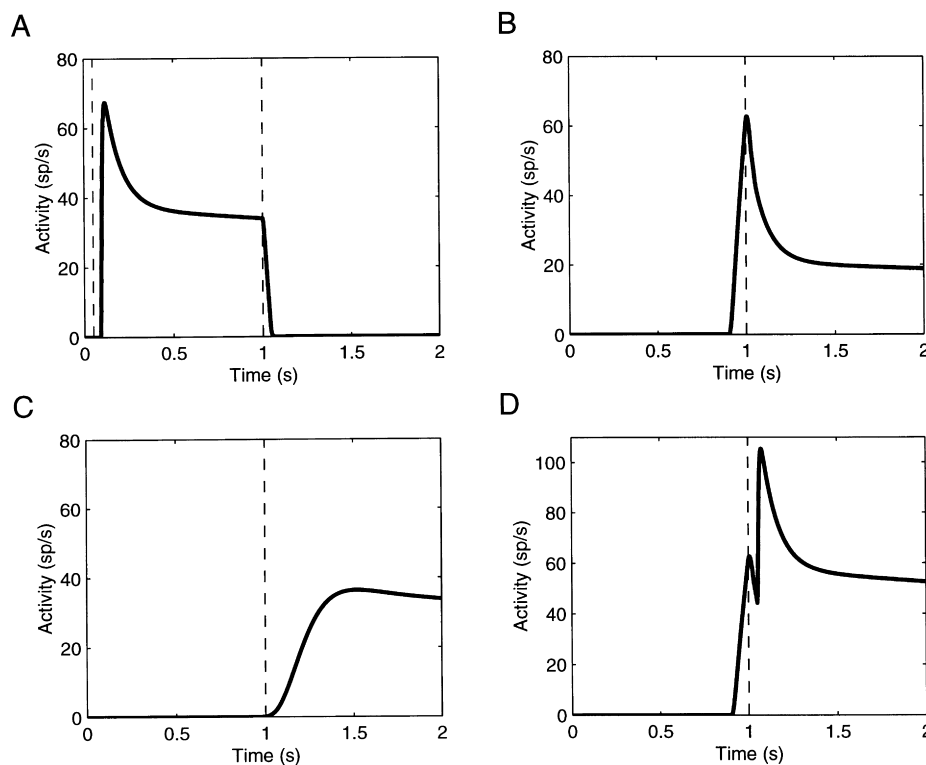


Fig. 13. Simulations of the model. At time 50 ms (left dashed vertical line in panel A) the target is flashed 30° to the right of the fovea. At time 1 s a 10° rightward saccade is executed. The behavior of LIP_{PT} cells is illustrated. (A) 30° LIP_{PT} cell. The cell bursts 50 ms after target onset, and then has a sustained activity. When the saccade starts, its activity starts decaying, and it is over by saccade end. (B) 20° LIP_{PT} cell with a predictive response. Approximately 90 ms before saccade onset the cell starts discharging, it bursts and then has a sustained activity. The remapping has occurred. (C) 20° LIP_{PT} cell that shows remapping with a longer latency. The cell starts discharging after saccade onset, has a small burst and then a sustained activity. (D) Same cell as in B, but the target is kept on for the duration of the trial. 60 ms after the end of the saccade there is a burst due to the reafferent visual response.

resemblance between Fig. 13D and Fig. 7B, which was obtained recording in LIP under the conditions simulated here.

5.3. Biological plausibility of the model

The model presented here requires a specific connectivity and some non-trivial operations performed at the dendritic level. The natural questions is: ‘Is this just another exercise in data fitting or is it somehow related to the way the brain is organized?’’. We now present some arguments for the model’s biological plausibility.

First of all, we think that there is no circuitry simpler than that illustrated in Fig. 10 that can account for the responses of LIP neurons to stable and flashed stimuli using only a step-like visual input. In fact, this circuitry accounts for the relationship between intensity of the phasic and intensity of the tonic response (Barash et al., 1991b), for the different cell types found in LIP, does not require the existence of undiscovered neurons, and it does not ignore critical functions (e.g. by saying that some other area memorizes the target location).

However, we have to admit that the precise pairings between LIP_{PT} cells and FEF_{MOV} cells at the level of the LIP_P cells illustrated in Fig. 11 appear at best suspicious. What makes this connectivity less unrealistic is that it can be learned very easily, using a local Hebbian rule to potentiate and depress the synapses while the subject looks around at natural scenes. In fact, suppose that initially all the possible pairings exist, with random and relatively weak connections to the LIP_P cells; now, if we potentiate the synapses on the LIP_P cell that discharge in conjunction with (or just before) the LIP_P cell, and we depress those that do not fire when the LIP_P cell fires, and then we look around at natural scenes (i.e. scenes where the targets are stable and not flashing) the only pairs that will survive will be those that we have proposed. Furthermore, we only hypothesize such connectivity for a very limited subset of neurons, i.e. those that show the earliest remapping.

The circuitry illustrated in Fig. 12 is necessary to reproduce the range of latencies of the remapping observed in all areas. It is undeniable that it requires a fairly specific connectivity, but we do not see many other solutions to the problem. In fact, even schemes that rely on eye position signals to perform the remapping (Dominey and Arbib, 1992; Salinas and Abbott, 1995; Zipser and Andersen, 1988) (and that in any case can not reproduce the predictive remapping), would most likely need a similar connectivity.

With these arguments we have shown that the complexity of the model presented here is simply due to the need to account for a complex data set, and that any models that use different solutions to the remapping problem will need to make similar assumption and to postulate such a specific connectivity to account for the body of data that our model replicates.

6. Discussion

We have reviewed evidence that some neurons in LIP, FEF and SC always encode a target location relative to the direction of gaze. Ordinarily they respond to targets that appear in a certain region of the retina. However, these neurons also respond (in some cases predictively) when a saccade brings a target into their receptive field. In other words these neurons always provide the saccadic system with a displacement signal that is appropriate to foveate a target.

Furthermore, we have presented a model of LIP–FEF interactions that reproduces the pattern of activity of neurons that shift their receptive fields without requiring an explicit representation of the location of targets in space. These results show that absolute target position is not necessary for generating a spatially accurate vector across eye movements and that predictive responses can easily arise, provided that signals relative to the impending motor behavior, and not just to the position of the eyes, are used. Further support for an oculocentric mapping of saccade targets has been provided using, instead of the neurophysiological and clinical methods described above, a clever behavioral test (Henriques et al., 1998).

Recently Gottlieb et al. (1998) have shown that the LIP neurons we have described and modeled here provide a map of the salient visual environment unrelated to the planning and performance of purposive saccades. These cells exhibit remapping under circumstances where the monkey could never intend a saccade to the stimulus. Accordingly, we think that the conclusions we have drawn here can be regarded as holding not just for maps involved in a specific sensori-motor transformation but for maps of saliency in general. From an evolutionary point of view, this makes perfect sense. In fact, if we have determined that a part of the visual scene is particularly meaningful, we do not want to lose track of it every time the eyes move. Thus, the ability to predict where on the retina an area of interest will end up after a movement could result in a behavioral advantage.

It is always possible that the areas that show predictive visual responses and spatially accurate visual memory responses generate them by accessing an explicit representation of target position in space that lies elsewhere. Zipser and Andersen (1988), for example, have pointed out that it is theoretically possible to calculate absolute target position in space from neurons whose visual responses, although retinotopic, are modulated by the position of the eye in the orbit. Such responses are found in many visual areas including the parietal cortex (Andersen and Mountcastle, 1983), V3A (Galletti and Battaglini, 1989) and the lateral geniculate nucleus (Lal et al., 1990); and one could argue that representing the location of targets in eye-centered coordinates is not very useful, because we might want to point to a target with our hand instead of redirecting our gaze to it.

However, regardless of the representation used to store the target location, the brain is still faced with the problem

of computing the proper coordinate transformation to produce a movement or more in general to guide behavior, and starting from absolute, spatial coordinates as opposed to eye-centered coordinates does not simplify dramatically such a transformation. The major difference is that to work in absolute coordinates one has to postulate the existence of neural signals that have not been reported (like target location in space or a predictive eye position signal); on the other hand, if the targets are memorized in continuously updated eye-centered coordinates, we have shown here that the signals that have been already reported in neurophysiological studies suffice.

Acknowledgements

C.Q. was partially supported by a grant (Sistemi naturali ed artificiali nei problemi cognitivi e dell'apprendimento) from Ministero dell'Università e della Ricerca Scientifica e Tecnologica to Paolo Inchingolo.

References

- Andersen R.A., & Mountcastle V.B. (1983). The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex. *Journal of Neuroscience*, 3, 532–548.
- Barash S., Bracewell R.M., Fogassi L., Gnadt J.W., & Andersen R.A. (1991). Saccade-related activity in the lateral intraparietal area. I. Temporal properties. *Journal of Neurophysiology*, 66, 1095–1108.
- Barash S., Bracewell R.M., Fogassi L., Gnadt J.W., & Andersen R.A. (1991). Saccade-related activity in the lateral intraparietal area. II. Spatial properties. *Journal of Neurophysiology*, 66, 1109–1124.
- Bruce C.J., & Goldberg M.E. (1985). Primate frontal eye fields: I. Single neurons discharging before saccades. *Journal of Neurophysiology*, 53, 603–635.
- Bushnell M.C., Goldberg M.E., & Robinson D.L. (1981). Behavioral enhancement of visual responses in monkey cerebral cortex: I. Modulation in posterior parietal cortex related to selective visual attention. *Journal of Neurophysiology*, 46, 755–772.
- Dominey P.F., & Arbib M.A. (1992). A cortico-subcortical model for generation of spatially accurate sequential saccades. *Cerebral Cortex*, 2, 153–175.
- Duhamel J.-R., Colby C.L., & Goldberg M.E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255, 90–92.
- Duhamel J.-R., Goldberg M.E., FitzGibbon E.J., Sirigu A., & Grafman J. (1992). Saccadic dysmetria in a patient with a right frontoparietal lesion: the importance of corollary discharge for accurate spatial behavior. *Brain*, 115, 1387–1402.
- Galletti C., & Battaglini P.P. (1989). Gaze-dependent visual neurons in area V3A of monkey prestriate cortex. *Journal of Neuroscience*, 9, 1112–1125.
- Goldberg M.E., & Bruce C.J. (1990). Primate frontal eye fields. III. Maintenance of a spatially accurate saccade signal. *Journal of Neurophysiology*, 64, 489–508.
- Goldberg M.E., & Bushnell M.C. (1981). Behavioral enhancement of visual responses in monkey cerebral cortex. II. Modulation in frontal eye fields specifically related to saccades. *Journal of Neurophysiology*, 46, 773–787.
- Gottlieb J., Kusunoki M., & Goldberg M.E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, 391, 481–484.
- Hallett P.E., & Lightstone A.D. (1976). Saccadic eye movements to flashed targets. *Vision Research*, 16, 107–114.
- Heide W., Blankenburg M., Zimmermann E., & Kompf D. (1995). Cortical control of double-step saccades – Implications for spatial orientation. *Annals of Neurology*, 38, 739–748.
- Henriques D.Y.P., Klier E.M., Smith M.A., Lowy D., & Crawford J.D. (1998). Gaze-centered remapping of remembered visual space in an open-loop pointing task. *Journal of Neuroscience*, 18, 1583–1594.
- Jürgens R., Becker W., & Kornhuber H.H. (1981). Natural and drug-induced variations of velocity and duration of human saccadic eye movements: evidence for a control of the neural pulse generator by local feedback. *Biological Cybernetics*, 39, 87–96.
- Kusunoki M., Colby C.L., & Goldberg M.E. (1994). Perisaccadic changes in the excitability of visual neurons in monkey parietal cortex. *Society for Neuroscience Abstract*, 20, 773.
- Lal R., Friedlander M.J., & Brunet P. (1990). Effect of passive eye position changes on retinogeniculate transmission in the cat. *Journal of Neurophysiology*, 63, 502–522.
- Mays L.E., & Sparks D.L. (1980). Dissociation of visual and saccade-related responses in superior colliculus neurons. *Journal of Neurophysiology*, 43, 207–232.
- Mays L.E., & Sparks D.L. (1980). Saccades are spatially, not retinocentrically, coded. *Science*, 208, 1163–1165.
- Paré M., & Wurtz R.H. (1997). Discharge properties of lateral intraparietal sulcus neurons antidromically activated by superior colliculus stimulation. *Society of Neuroscience Abstracts*, 23, 843.
- Robinson, D.A. (1975). Oculomotor control signals. In G. Lennerstrand & P. Bach-y-Rita (Eds.), *Basic mechanisms of ocular motility and their clinical implications* (pp. 337–374). Oxford: Pergamon Press.
- Salinas E., & Abbott L.F. (1995). Transfer of coded information from sensory to motor networks. *Journal of Neuroscience*, 15, 6461–6474.
- Schlag J., & Schlag-Rey M. (1991). Colliding saccades may reveal the secret of their marching orders. *Trends in Neuroscience*, 13, 410–415.
- Umeno M.M., & Goldberg M.E. (1997). Spatial processing in the monkey frontal eye field. I. Predictive visual responses. *Journal of Neurophysiology*, 78, 1373–1383.
- Walker M.F., FitzGibbon E.J., & Goldberg M.E. (1995). Neurons in the monkey superior colliculus predict the result of impending saccadic eye movements. *Journal of Neurophysiology*, 73, 1988–2003.
- Westheimer, G. (1989). History and methodology. In R.H. Wurtz & M.E. Goldberg (Eds.), *The neurobiology of saccadic eye movements*, Reviews of Oculomotor Research, Vol. III (pp. 3–12). Amsterdam: Elsevier.
- Zipser D., & Andersen R.A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331, 679–684.