Spatial constancy and the brain: insights from neural networks

Robert L. White III and Lawrence H. Snyder*

*Department of Anatomy and Neurobiology, Washington University School of Medicine,
Box 8108, 660 South Euclid Avenue, St Louis, MO 63110, USA*

To form an accurate internal representation of visual space, the brain must accurately account for movements of the eyes, head or body. Updating of internal representations in response to these movements is especially important when remembering spatial information, such as the location of an object, since the brain must rely on non-visual extra-retinal signals to compensate for self-generated movements. We investigated the computations underlying spatial updating by constructing a recurrent neural network model to store and update a spatial location based on a gaze shift signal, and to do so flexibly based on a contextual cue. We observed a striking similarity between the patterns of behaviour produced by the model and monkeys trained to perform the same task, as well as between the hidden units of the model and neurons in the lateral intraparietal area (LIP). In this report, we describe the similarities between the model and single unit physiology to illustrate the usefulness of neural networks as a tool for understanding specific computations performed by the brain.

Keywords: reference frames; coordinate frames; eye movements; monkey; cortex

1. INTRODUCTION

As you read this article, your eyes make a sequence of rapid movements (saccades) to direct your gaze to the words on this page. Although saccades move the eyes at a speed of up to 400° per second, you do not perceive that the visual world is moving rapidly. The characters, which are stationary on the page, do not appear to move, even though their projection onto the retina changes. The ability of the visual system to account for self-motion is called space constancy (Holst & Mittelstaedt 1950; Helmholtz 1962; Stark & Bridgeman 1983; Bridgeman 1995; Deubel *et al.* 1998) and is an important component of goal-directed behaviour. In order to interact with objects in our surroundings, we must account for motion of the eyes, head and body for accurate visually guided movements (e.g. saccades, reaching).

Space constancy is particularly important when remembering the locations of objects that become occluded from view. In this case, the compensation for self-motion is critical for maintaining an accurate representation of an object's location while it is being remembered. Both humans and monkeys are capable of performing such compensation, which is also known as spatial updating.

To update accurately, the brain must synthesize information regarding movements of various types (eye-in-head, head-on-body and body-in-world) with remembered visuospatial information. However, multiple possible solutions exist. It has been suggested

that spatial locations could be stored with respect to an allocentric, world-fixed coordinate system that does not vary with the motion of the observer. An object that is stationary in the world would be stationary in such a world-fixed coordinate frame. However, not all objects are stationary in the world, and often their motion is linked to movements of the observer. For example, when tracking a moving object, features of the moving object remain stationary with respect to the retina. Similarly, objects fixed with respect to the head or body move on the retina, but remain stationary in a head- or body-fixed reference frame. Thus, one would need a multitude of redundant static representations of space to encode the variety of ways in which the objects move in the world.

There is evidence that the primate brain contains multiple spatial representations in different coordinate frames (e.g. Galletti *et al.* 1993; Duhamel *et al.* 1997; Snyder *et al.* 1998; Olson 2003; Martinez-Trujillo *et al.* 2004; Park *et al.* 2005). However, a large number of representations are sensoritopic, that is, represented with respect to the spatial characteristics of the input. Visual-spatial representations are often gaze-centred, auditory-spatial representations are often head-centred, etc. How might a sensoritopic neural representation update spatial locations in response to shifts in gaze? In the lateral intraparietal area (LIP), a cortical area in monkeys known to be involved in the planning of saccadic eye movements (Andersen *et al.* 1992), neurons have gaze-centred receptive fields (Colby *et al.* 1995), but the responsiveness (gain) to targets presented in the receptive field can be linearly modulated by changes in gaze position. This modulation is called a gaze position gain field (Andersen & Mountcastle 1983).

Neural network modelling has been a useful tool in elucidating the mechanisms of spatial updating.

* Author for correspondence (larry@eye-hand.wustl.edu).

Electronic supplementary material is available at <http://dx.doi.org/10.1098/rstb.2006.1965> or via <http://www.journals.royalsoc.ac.uk>.

One contribution of 15 to a Theme Issue 'The use of artificial neural networks to study perception in animals'.

Zipser & Andersen (1988) demonstrated gaze position gain fields in hidden layer units that combine retinal and gaze position information in a feedforward neural network. Xing & Andersen (2000) extended these results and showed that a recurrent neural network could dynamically update a stored spatial location based on changes in a position signal. In this network, gain fields for position were observed in the hidden layer. It has been proposed that gain fields might serve as a general mechanism for updating in response to changes in gaze, since gain fields have been observed in many of the areas where spatial signals are updated (Andersen *et al.* 1990).

Gaze position signals are not the only means by which to update a remembered location in a gaze-centred coordinate system. Droulez & Berthoz (1991) demonstrated that world-fixed targets could be updated entirely within a gaze-centred coordinate system, with the use of a gaze velocity signal. The use of velocity signals may be particularly relevant for updating with vestibular cues (signals from the inner ear that sense rotation and translation of the head). For example, humans and monkeys are able to update accurately in response to whole-body rotations, and neurons in LIP appear to reflect this updating (Baker *et al.* 2002). In the monkey, gain fields for body position have not been observed in LIP (Snyder *et al.* 1998). Perhaps, vestibular velocity signals provide the necessary inputs for updating, in this case without a necessity for a gain field mechanism.

In order to explore how spatial memories are stored and manipulated in complex systems, we compared the responses of macaque monkeys and individual LIP neurons with the responses of recurrent neural networks and individual network nodes. We trained two monkeys to memorize the location of a visual target in either a world- or gaze-fixed frame of reference. The reference frame was indicated by the colour of the fixation point, as well as an instructional sequence at the beginning of each trial. After the target disappeared, the animals' gaze was shifted horizontally to a new location by a smooth pursuit eye movement, a saccade, or a whole-body rotation. In whole-body rotations, the fixation point moved with the same angular velocity as the monkey so that the gaze remained aligned with the body. Following the gaze shift, the fixation point disappeared, cuing the animal to make a memory-guided saccade to the remembered location of the target. Figure 1 shows a schematic of the task. Subsequent analyses in this report focus on the whole-body rotation case.

Next, we created and trained a recurrent neural network model to perform an analogue of the animals' task. The network consisted of an input layer, a fully recurrent hidden layer and an output layer. The input nodes had activations corresponding to the retinal input, gaze position and gaze velocity, and the reference frame to be used. The responses of the output nodes were decoded in order to determine the saccadic goal location in gaze-centred coordinates. The model only traced one dimension of visual space (analogous to the horizontal component of the monkeys' task). We first established that our network was a reasonable model of the animals' behaviour, and then used the network

Figure 1. Schematic of the task. Animals began each trial by looking at a fixation point. The target was briefly flashed, and could appear at one of the many different locations. It disappeared, and gaze was shifted either to the left or the right, either by whole-body rotation (VOR cancellation), smooth pursuit or a saccade. The delay period concluded with the disappearance of the fixation point. The animal was rewarded for moving either to the world- or the gaze-fixed location of the original target. The neural network model tracked only the horizontal component of target location. The network was given a transient visual target and a sustained reference frame cue. A slow gaze shift followed, after which the network was queried for a saccade vector to either the world- or the gaze-fixed location of the original target.

model to investigate issues that are difficult to investigate in the animal. We were principally interested in the inputs used to drive updating and whether or not the mechanism for updating involves gain fields.

2. MATERIAL AND METHODS

We implemented a three-layered recurrent neural network model (White & Snyder 2004a), trained it on a flexible updating paradigm, and compared the output and internal activity of the model with the behaviour and neuronal activity obtained from monkeys performing an analogous task (Baker *et al.* 2003). Detailed methods are available online as electronic supplementary material.

The model's input layer was designed to represent a visual map in one dimension of space and to provide information about gaze position and gaze velocity. Output was in eye-centred coordinates, representing the goal location of an upcoming saccade. For flexible behaviour, we added a binary cue that indicated whether the target was world- or gaze-fixed. The network had full feedforward connections between successive layers and full recurrent connections in the hidden layer. The recurrent connections provided the network with memory over time.

The task of the network was to store and, if necessary, transform a pattern of activity representing the spatial location of a target into a motor command. In order to correctly perform the task, the network had to either compensate or ignore any changes in gaze that occurred during the storage interval, depending on the instruction provided by the two reference frame input units. The network output provides an eye-centred representation of the stored spatial location. The network was trained using a 'backpropagation through time' algorithm (Rumelhart *et al.* 1986).

Figure 2. (a) Accuracy of animal behaviour (i) and network output (ii). (i) Mean horizontal saccade amplitudes versus ideal saccade amplitudes from two monkeys (Baker *et al.* 2003). (ii) Mean coded output ($\pm$ s.e.) from three networks of 35 hidden units at the last time step versus the ideal saccade amplitude for world-fixed (grey) and subject-fixed (black) trials. The straight black line in each panel indicates perfect performance (actual = ideal). (b) Precision of animal saccades and network output, measured as the standard deviation (s.d.) of the saccade amplitudes, for the same data as shown in (a).

In an analogous task, two rhesus monkeys performed memory-guided saccades to world- and gaze-fixed target locations following whole-body rotations (vestibulo-ocular reflex (VOR) cancellation), smooth pursuit eye movements or saccades that occurred during the memory period (Baker *et al.* 2003). Fixation point offset at the conclusion of the memory period cued monkeys to make a saccade to the remembered target location. Changes in eye position as well as the neuronal activity of single neurons in cortical area LIP were recorded.

3. RESULTS

Animals made accurate memory-guided saccades to target locations in the appropriate reference frame. Figure 2a(i) compares the average memory-guided saccade horizontal amplitude with the ideal amplitude. Animals were equally accurate in both reference frame conditions across a range of target locations, although animals tended to undershoot the most peripheral targets in both conditions. When we examined the precision of memory-guided saccade endpoints, we found a difference between the performance on world- and gaze-fixed trials. Figure 2b(i) shows the standard deviation of saccade endpoints as a function of ideal saccade amplitude. World-fixed saccades had a larger standard deviation than gaze-fixed saccades.

We found that our recurrent neural network model, once trained, had similar behavioural performance as the animals. We used the centre of mass of the output activations to 'read out' the saccadic goal location. When noise was injected into the inputs to produce variability in the model's output, the saccade readout was largely accurate (figure 2a(ii)), but precision was worse for world-fixed trials compared with gaze-fixed trials (figure 2b(ii)). This emergent property of the model's behaviour was not explicitly designed or trained into the architecture, and therefore indicates that our computational model captures at least one feature of spatial updating that occurs in biology.

We next examined the responses of single neurons in LIP during the flexible updating task. We found that many cells updated their activity in response to gaze shifts during world-fixed trials, but held their activity constant during gaze-fixed trials. This pattern of activity is consistent with dynamic updating of spatial memories in a gaze-centred coordinate system. The responses from a single cell are shown in figure 3a. Consider first the response to target presentation prior to the gaze shift figure 3a(i). A target appearing inside the cell's receptive field (RF) produced a transient burst of activity followed by sustained elevation in activity that continued after the target had disappeared from view (solid traces). A target appearing outside the cell's RF produced no change in activity, or even a slight decrease to below baseline (dotted traces).

Next, we considered the effect of a gaze shift. In figure 3a(ii), whole-body rotations ($\pm 20^\circ$; VOR cancellation paradigm) brought the world-fixed location of the target either into (solid traces) or out of (dotted traces) the cell's RF. A short time later (figure 3a(iii)), the animal was rewarded for an accurate memory-guided saccade that treated the target as either world-fixed (grey traces) or gaze-fixed (black traces). A world-fixed target whose remembered location moved from inside to outside the RF resulted in a decrease in firing rate (solid grey trace). Conversely, when the location moved from outside to inside the RF, the firing rate increased (dotted grey trace). Activity was unaffected by either type of gaze shift when the target was gaze-fixed (solid and dotted black traces).

Figure 3b shows the normalized average response from 42 LIP cells in one monkey during whole-body rotations of the same type as figure 3a. Like the example cell, it shows a pattern consistent with dynamic updating in a gaze-centred coordinate system. The population average activity decreases when a target's remembered location leaves the RF on world-fixed trials (solid grey trace), but does not change on gaze-fixed trials (solid black trace). Similarly, the population

Figure 3. (a) Smoothed instantaneous firing rates for a single LIP neuron. Mean firing rates (thick lines) $\pm$ s.e.m. (thin lines) are shown for targets presenting inside the neuron's RF (solid lines) and outside the RF (dashed lines). On world-fixed trials (grey), a subsequent 20° whole-body rotation to the left or the right brought the goal location out of or into the neuron's RF, resulting in a decrement or increment in the neuron's activity, respectively. On gaze-fixed trials (black), target locations and rotations were identical, but the goal location remained constant (either inside or outside the RF). Little modulation of activity was observed for this neuron during the delay on the gaze-fixed trials. Data are aligned on target presentation (i), middle of the rotation (ii) and the memory-guided saccade (iii). Seven trials were performed for each condition. On the bottom, the average horizontal eye position is shown for one of the conditions. (b) Smoothed instantaneous firing rates for the average activity from 42 LIP neurons during whole-body rotations. Same format as (a), except standard error traces are omitted for clarity.

activity increases when a target's remembered location enters the RF (dashed grey trace), but does not change during the gaze shift (dashed black trace).

It appears that on average, LIP cells change their firing rate (update) in response to shifts of gaze for world-fixed targets but not gaze-fixed targets. However, this updating is incomplete: in figure 3b, the difference between in- and out-RF activities is smaller on world-fixed trials than gaze-fixed trials. We quantify this below.

For an analogue of LIP cells in our model, we examined the responses of hidden (middle) layer units in the recurrent neural network. We found that many of the hidden units were active during the delay period and dynamically updated their activity on world-fixed trials, much like the LIP cells in the monkey. Activity from a single model unit is shown in figure 4a. The model unit responds much in the same way as the example cell from LIP (figure 3a). On gaze-fixed trials, in- and out-RF activities remain constant over the memory period, whereas, on world-fixed trials, the in-RF activity decreases and the out-RF activity increases following the gaze shift.

Figure 4b shows the population average of all 35 hidden units from the neural network. The patterns of neural activity are remarkably similar to those seen in the example cell and in the responses of LIP neurons.

Like LIP, it also appears that updating in the model is incomplete: the difference between in- and out-RF activities at the end of the trial is greater for gaze-fixed targets than for world-fixed targets.

We next wished to examine the relative amounts of gaze-fixed and world-fixed information conveyed by individual cells. In other words, we asked if individual cells were 'specialized' for one reference frame or the other, or 'flexible,' like the example cell (figure 4a). To address this question, we calculated the relative amounts of gaze- and world-fixed information conveyed by each unit using the receiver operating characteristic (ROC) of an ideal observer. We used the ROC to discriminate whether a target was inside or outside of a cell's RF, and calculated the cell's performance as the area under the ROC curve. If the individual cells were specialized for either the world- or gaze-fixed transformation, then we would expect individual cells to have ROC values that were greater than chance for one type of transformation, and at chance or systematically worse than chance for the other. In contrast, if the cells were capable of contributing to both transformations, their corresponding ROC values would be greater than their chance for both transformations.

Figure 5a displays the world- and gaze-fixed ROC discrimination values (the area under the ROC curve)

Figure 4. (a) Activity of one hidden layer unit during individual trials for a 20° gaze shift. When presented with a gaze-fixed cue and target inside its RF (solid black trace) or outside the RF (dashed black trace), there is little change in the unit's response over time. When a world-fixed target appears inside the unit's RF, the subsequent gaze shift moves the remembered target location outside the RF and the unit's activity decreases (solid grey trace). When a world-fixed target appears outside the unit's RF, but the gaze shift brings the remembered target location into the RF, the unit increases its activity (dashed grey trace). Grey shading indicates the time-steps over which the gaze shift occurred. (b) Average response of all 35 hidden units in a single trained network. Same format as (a). Error bars represent standard errors from three networks.

on rotation trials for 86 LIP cells from two monkeys. Each point represents one cell. There are many points in the upper right quadrant, indicative of many individual cells that contribute to both gaze- and world-fixed transformations. A substantial subset of cells lies at or below an ordinate value of 0.5 (lower right quadrant), indicating that they retain spatial information on gaze-fixed trials but either contribute nothing or even provide erroneous information on world-fixed trials. Very few neurons lie in the upper left quadrant (valid information on world-fixed trials and invalid information on gaze-fixed trials). The net result is that mean discrimination was better on gaze-fixed trials compared with world-fixed trials (mean ROC: 0.83 versus 0.58; $p < 1 \times 10^{-12}$, Wilcoxon test, $n = 86$).

We performed the same ROC analysis on the hidden units of the model (with injected noise to produce variability in the output). We found that units discriminated well in both types of trials (figure 5b), but that mean discrimination was better on gaze-fixed

Figure 5. (a) Comparison of gaze-fixed versus world-fixed discrimination for LIP neurons recorded during trials with whole-body rotation gaze shifts. For each unit, the area under the ROC curve (AUC) was calculated (see §2). Each unit is plotted with respect to its AUC for gaze-fixed (abscissa) and world-fixed (ordinate) trials. Perfect discrimination corresponds to an AUC of 1. Dashed lines indicate the level of chance performance (AUC of 0.5). (b) Same format as (a) for units in the hidden layer of the recurrent network model ($n = 35$).

trials than world-fixed trials (mean ROC: 0.90 versus 0.61; $p < 1 \times 10^{-5}$, Wilcoxon test, $n = 35$). This pattern mimics that seen for the LIP neurons following the whole-body rotations (figure 5a).

We next asked whether networks trained with both position and velocity signals relied preferentially on one of these inputs. To answer this question, we zeroed the gaze shift inputs of either the position or velocity input nodes of networks trained with both position and velocity inputs. This is analogous to selectively 'lesioning' either the position or velocity inputs to the network. The performance deficits that resulted indicated that these networks preferentially relied on velocity to perform spatial updating (figure 6a). It was not the case that a velocity signal was required for updating, because networks trained with either velocity or position alone performed quite well. Thus, it appears that models show a preference for velocity inputs for spatial updating.

In our last model experiment, we examined the presence of gaze position gain fields in networks trained

Figure 6. (a) Performance of the network when gaze signals are removed. RMS error for all world-fixed (grey) and gaze-fixed (black) trial types is shown for the network with position inputs removed, the intact network and the network with velocity inputs removed. The network has a strong preference for the velocity input over the position input. Error bars represent standard errors from three networks. (b) Proportion of hidden layer units demonstrating gaze position gain fields stronger than 0.2% per degree under the different gaze input conditions. Error bars represent standard errors from three networks.

with both gaze position and velocity inputs, position alone or velocity alone. Gaze position gain fields have been postulated to be involved in the mechanisms of spatial updating (Xing & Andersen 2000; Balan & Ferrera 2003). However, we postulated that hidden units of networks that relied on a gaze velocity signal would not demonstrate gaze position gain fields. When networks were trained with only gaze position signals, we found an abundance on hidden units with gain field responses. Inversely, when gaze velocity signals were used for training, virtually no gain fields were present in hidden layer units (figure 6b). Our results indicate that gaze position gain fields are not a necessary feature of networks that perform spatial updating.

4. DISCUSSION

We have described a three-layer recurrent neural network model that combines retinal and extra-retinal signals, and is capable of encoding target locations as either fixed in the world or fixed with respect to gaze. We found the output of the network mimics the pattern of behaviour in animals trained to perform a similar task. Furthermore, the response properties of the

hidden (middle) layer units in this network resemble those of neurons in LIP. By removing specific inputs to the model, we tested whether updating preferentially utilizes gaze position or gaze velocity signals, and found that the network strongly preferred velocity for updating world-fixed targets. We found that gaze position gain fields were present when position signals, but not velocity signals, were available for updating.

Neural networks provide a powerful analytic tool in understanding computations performed by the brain. They provide a means of analysing and dissecting distributed coding and computation by a population of neurons. Furthermore, they provide a means to test insightful manipulations that would be difficult to perform *in vivo*, such as the input lesion experiments described above.

One of the goals of our modelling experiments was to address whether a single population of neurons could perform both world- and gaze-fixed transformations dynamically. We found that a neural network model could perform both world- and gaze-fixed transformations in a flexible manner. Furthermore, the precision of memory-guided behaviour was worse on world-fixed trials compared with gaze-fixed trials, consistent with dynamic updating in a gaze-centred coding scheme (Baker *et al.* 2003). Monkey behaviour showed an identical pattern, suggesting that a gaze-centred coding scheme could be utilized to track remembered target locations in the brain.

LIP has a putative role in spatial working memory (Gnadt & Andersen 1988). Neurons in this area encode spatial information using gaze-centred coordinates (Colby *et al.* 1995). When we examined the responses of single LIP neurons during a spatial updating task, we found they were markedly similar to those in the hidden layer of the model. Both neurons and model units had gaze-centred RFs and, on an average, updated their delay-period activity during world-fixed trials and (correctly) maintained their activity during gaze-fixed trials. The marked similarity between LIP cells and model units suggests that LIP may contribute to the flexible updating in the brain necessary for correct behavioural output. However, it is quite possible that this neural activity is merely reflective of computations that occur elsewhere in the brain—a limitation of many electrophysiological recording experiments.

The second goal of the modelling experiments was to ask whether world- and gaze-fixed spatial memories were held by sub-populations of units that specialized in just one of the two transformations, or by a single homogeneous population of units that flexibly performed both transformations. In the model, the world- and the gaze-fixed transformations were not performed by separate sub-networks, but by a distributed population of units that contributed to both world- and gaze-fixed output. Using the same analysis, we similarly found that cells in LIP did not specialize in either the world- or the gaze-fixed transformation.

Spatial updating requires an estimate of the amplitude and direction of movements that have been made while a location is maintained in memory. Such an estimate could arise from sensory signals or from

motor commands. By utilizing a copy of a motor command, areas involved in spatial updating could quickly access gaze shift signals, without the delay associated with sensory feedback (e.g. proprioception). An input carrying such a signal would be, in essence, a duplicate of the motor command itself, and has been called either a 'corollary discharge' or 'efference copy' signal. Recent studies have identified a role for saccadic motor commands in spatial updating following shifts of eye position (Guthrie *et al.* 1983; Sommer & Wurtz 2002). New evidence points to frontal eye field (FEF) as a possible locus of updating in this case (White & Snyder 2004b; Opris *et al.* 2005).

If corollary discharge signals are used for updating, what is the nature of these signals? Do they report gaze position or gaze velocity? We examined how both gaze position and gaze velocity signals are used in spatial updating by testing our recurrent network model. First, we found that in networks trained with both gaze velocity and gaze position signals, removing the velocity input produced a profound deficit in performance, whereas removing the position inputs produced only a minor effect (figure 6a). In the model, velocity was preferentially utilized to perform the updating computation while position only played a minor role. Second, we looked for hidden units with position-dependent gain field responses in networks trained with only one of the signals or with both signals in combination. We found that networks trained with gaze velocity did not demonstrate gaze position gain fields (figure 6b). By contrast, networks trained with gaze position contained many units with gain fields.

Recent models have highlighted the importance of a gaze position gain field representation in updating for changes in gaze (Cassanello & Ferrera 2004; Smith & Crawford 2005). Results from our model would indicate that such a representation would rely on gaze position signals (figure 6b). However, we suggest that although position signals may be used, they are not required for spatial updating. As an example, we have shown that LIP neurons update in response to whole-body rotations (Baker *et al.* 2002) but lack gain fields for body position (Snyder *et al.* 1998). Based on results from our recurrent network model (figure 6b), the LIP neurons could update in response to whole-body rotations if gaze shift signals were provided by vestibular velocity signals. Updating with velocity signals would not require that positional gain fields be present in LIP.

Based on behavioural differences, spatial updating for saccades could be a special case. In a recent study from our laboratory, a distinct pattern of updating behaviour was observed following saccades (as compared to rotations and pursuit; Baker *et al.* 2003). Updating for saccades might use signals conveying eye velocity (Droulez & Berthoz 1991) or eye displacement (White & Snyder 2004a). However, results from our model (figure 6b) predict that such a system would not demonstrate gain fields for eye position. Since eye position gain fields are observed in LIP and in FEF (Andersen *et al.* 1990; Balan & Ferrera 2003), we support the hypothesis that saccadic updating relies on a gain field mechanism.

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. *J. Neurosci.* **3**, 532–548.
- Andersen, R. A., Bracewell, R. M., Barash, S., Gnadt, J. W. & Fogassi, L. 1990 Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J. Neurosci.* **10**, 1176–1196.
- Andersen, R. A., Brotchie, P. R. & Mazzoni, P. 1992 Evidence for the lateral intraparietal area as the parietal eye field. *Curr. Opin. Neurobiol.* **2**, 840–846. (doi:10.1016/0959-4388(92)90143-9)
- Baker, J. T., White, R. L. & Snyder, L. H. 2002 Reference frames and spatial memory operations: area LIP and saccade behavior. In *Society for neuroscience*, Program No. 57.16.
- Baker, J. T., Harper, T. M. & Snyder, L. H. 2003 Spatial memory following shifts of gaze. I. Saccades to memorized world-fixed and gaze-fixed targets. *J. Neurophysiol.* **89**, 2564–2576. (doi:10.1152/jn.00610.2002)
- Balan, P. F. & Ferrera, V. P. 2003 Effects of gaze shifts on maintenance of spatial memory in macaque frontal eye field. *J. Neurosci.* **23**, 5446–5454.
- Bridgeman, B. 1995 A review of the role of efference copy in sensory and oculomotor control systems. *Ann. Biomed. Eng.* **23**, 409–422.
- Cassanello, C. R. & Ferrera, V. P. 2004 Vector subtraction using gain fields in the frontal eye fields of macaque monkeys. In *Society for neuroscience*, Program No. 186.11.
- Colby, C. L., Duhamel, J. R. & Goldberg, M. E. 1995 Oculocentric spatial representation in parietal cortex. *Cereb. Cortex* **5**, 470–481.
- Deubel, H., Bridgeman, B. & Schneider, W. X. 1998 Immediate post-saccadic information mediates space constancy. *Vis. Res.* **38**, 3147–3159. (doi:10.1016/S0042-6989(98)00048-0)
- Droulez, J. & Berthoz, A. 1991 A neural network model of sensoritopic maps with predictive short-term memory properties. *Proc. Natl Acad. Sci. USA* **88**, 9653–9657. (doi:10.1073/pnas.88.21.9653)
- Duhamel, J. R., Bremmer, F., BenHamed, S. & Graf, W. 1997 Spatial invariance of visual receptive fields in parietal cortex neurons. *Nature* **389**, 845–848. (doi:10.1038/39865)
- Galletti, C., Battaglini, P. P. & Fattori, P. 1993 Parietal neurons encoding spatial locations in craniotopic coordinates. *Exp. Brain Res.* **96**, 221–229. (doi:10.1007/BF00227102)
- Gnadt, J. W. & Andersen, R. A. 1988 Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.* **70**, 216–220.
- Guthrie, B. L., Porter, J. D. & Sparks, D. L. 1983 Corollary discharge provides accurate eye position information to the oculomotor system. *Science* **221**, 1193–1195. (doi:10.1126/science.6612334)
- Helmholtz, H. V. 1962 *A treatise on physiological optics*. New York, NY: Dover.
- von Holst, E. & Mittelstaedt, H. 1950 The reafferent principle: reciprocal effects between central nervous system and periphery. *Naturwissenschaften* **37**, 464–476. (doi:10.1007/BF00622503)
- Martinez-Trujillo, J. C., Medendorp, W. P., Wang, H. & Crawford, J. D. 2004 Frames of reference for eye-head gaze commands in primate supplementary eye fields. *Neuron* **44**, 1057–1066. (doi:10.1016/j.neuron.2004.12.004)
- Olson, C. R. 2003 Brain representation of object-centered space in monkeys and humans. *Annu. Rev. Neurosci.* **26**, 331–354. (doi:10.1146/annurev.neuro.26.041002.131405)

- Opris, I., Barborica, A. & Ferrera, V. P. 2005 Effects of electrical microstimulation in monkey frontal eye field on saccades to remembered targets. *Vis. Res.* **45**, 3414–3429. (doi:10.1016/j.visres.2005.03.014)
- Park, J., Schlag-Rey, M. & Schlag, J. 2005 Frames of reference for saccadic command, tested by saccade collision in the supplementary eye field. *J. Neurophysiol.* **95**, 159–170. (doi:10.1152/jn.00268.2005)
- Rumelhart, D. E., Hinton, G. E. & Williams, R. 1986 Learning internal representations by error propagation. In *Parallel distributed processing: explorations in the microstructure of cognition* (eds D. E. Rumelhart & J. L. McClelland), pp. 316–362. Cambridge, MA: MIT Press.
- Smith, M. A. & Crawford, J. D. 2005 Distributed population mechanism for the 3-D oculomotor reference frame transformation. *J. Neurophysiol.* **93**, 1742–1761. (doi:10.1152/jn.00306.2004)
- Snyder, L. H., Grieve, K. L., Brotchie, P. & Andersen, R. A. 1998 Separate body- and world-referenced representations of visual space in parietal cortex. *Nature* **394**, 887–891. (doi:10.1038/29777)
- Sommer, M. A. & Wurtz, R. H. 2002 A pathway in primate brain for internal monitoring of movements. *Science* **296**, 1480–1482. (doi:10.1126/science.1069590)
- Stark, L. & Bridgeman, B. 1983 Role of corollary discharge in space constancy. *Percept. Psychophys.* **34**, 371–380.
- White 3rd, R. L. & Snyder, L. H. 2004a A neural network model of flexible spatial updating. *J. Neurophysiol.* **91**, 1608–1619. (doi:10.1152/jn.00277.2003)
- White, R. L. & Snyder, L. H. 2004 Delay period microstimulation in the frontal eye fields updates spatial memories. In *Society for neuroscience*, Program No. 527.9.
- Xing, J. & Andersen, R. A. 2000 Memory activity of LIP neurons for sequential eye movements simulated with neural networks. *J. Neurophysiol.* **84**, 651–665.
- 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. (doi:10.1038/331679a0)