Journal of Computational Neuroscience 10, 127–150, 2001

Self-Organizing Task Modules and Explicit Coordinate Systems in a Neural Network Model for 3-D Saccades

MICHAEL A. SMITH AND J. DOUGLAS CRAWFORD
Centre for Vision Research; Medical Research Council Group for Action and Perception; and Department of Psychology, York University, Toronto, Ontario, Canada, M3J 1P3
mas@yorku.ca

Received February 18, 2000; Revised August 22, 2000; Accepted September 15, 2000

Action Editor: B. Richmond

Abstract. The goal of this study was to train an artificial neural network to generate accurate saccades in Listing's plane and then determine how the hidden units performed the visuomotor transformation. A three-layer neural network was successfully trained, using back-prop, to take in oculocentric retinal error vectors and three-dimensional eye orientation and to generate the correct head-centric motor error vector within Listing's plane. Analysis of the hidden layer of trained networks showed that explicit representations of desired target direction and eye orientation were not employed. Instead, the hidden-layer units consistently divided themselves into four parallel modules: a dominant "vector-propagation" class (∼50% of units) with similar visual and motor tuning but negligible position sensitivity and three classes with specific spatial relations between position, visual, and motor tuning. Surprisingly, the vector-propagation units, and only these, formed a highly precise and consistent orthogonal coordinate system aligned with Listing's plane. Selective "lesions" confirmed that the vector-propagation module provided the main drive for saccade magnitude and direction, whereas a balance between activity in the other modules was required for the correct eye-position modulation. Thus, contrary to popular expectation, error-driven learning in itself was sufficient to produce a "neural" algorithm with discrete functional modules and explicit coordinate systems, much like those observed in the real saccade generator.

Keywords: saccades, visuomotor, coordinate systems, neural networks

1. Introduction

In the course of generating accurate visually guided behavior, the brain must transform eye-centered sensory signals into commands for movements relative to the head or body (Snyder et al., 1998; Soechting and Flanders, 1989; Batista et al., 1999; Flanders et al., 1999; Colby et al., 1995; Goldberg and Bruce, 1990). Most models of the visuomotor transformation address this *reference frame* problem by developing head-centric visual representations through comparisons with eye orientation (Zee et al., 1976; Grossberg and Kuperstein, 1986; Zipser and Andersen, 1988). Alternatively, it has been argued that the problem disappears if visual signals are mapped onto motor commands as a sequence of vector-displacement codes (Woodwoth, 1899), particularly for simple movements like saccades (Jürgens et al., 1981; Raphan, 1998). For example, according one popular model for the generation of saccades, remembered visual targets are remapped relative to current gaze direction during each saccade, so that subsequent saccades can be made without any internal reference to eye position (Moschovakis and Highstein, 1994; Colby and Goldberg, 1999).

However, when the actual three-dimensional (3-D) geometry of saccades is considered, vector displacement signals are *not* frame-independent—that is, the sensory vector is oculocentric, whereas the motor vector is headcentric. Thus, a fixed visuomotor mapping would produce inaccurate saccades, depending on initial eye position (Crawford and Guitton, 1997). Since the real saccade generator does not do this (Klier and Crawford, 1998), it follows that the saccade generator must be performing the correct position-dependent transformation. This does not necessarily contradict the idea of remapping visual targets in retinal coordinates (Colby and Goldberg, 1999), but it does affirm that such representations—or at least those chosen for motor execution—must then be put through an eye-to-head reference frame transformation at some point downstream (Henriques et al., 1998).

This being the case, the neuroscientist would then want to know where and how this transformation occurs in the brain. But at this time one is hampered in answering this question by the lack of any clear notion of what to look for in the neural signals. The "black-box" aspects of these transformations can be modeled with the use of operations like quaternion multiplication and "desired eye orientation" commands (Crawford and Guitton, 1997), but these are not likely to accurately represent the detailed operations of neural networks (Robinson, 1992). One useful approach, then, might be to first train a neural network to perform such tasks and determine how it solves the problem before tackling the physiological system directly.

Unfortunately, analysis of artificial neural networks has often proven just as difficult as analysis of real neural networks. One apparent success story germane to the current topic was the discovery that homogeneously distributed position-dependent "gain fields" on a retinotopic map can be used to construct either a map of space in headcentric coordinates or motor commands of the type required for 3-D saccades (Zipser and Andersen, 1988; Van Opstal et al., 1995; Liu et al., 1997). One criticism of this approach is that it still requires output signals (such as, headcentric spatial maps) that are physiologically unrealistic (Moschovakis and Highstein, 1994; Colby and Goldberg, 1999). But in theory, such codes may not be necessary (Zipser and Andersen, 1988). For example, it seems plausible that a neural network could use eye-position signals to transform retinally coded visual displacement signals into headcentric saccade displacement signals—as in Crawford and Guitton (1997)—without such intermediaries.

If so, then how would the network do it? Potentially, it could use an algorithm similar to that used in the "black-box" models (Crawford and Guitton, 1997; Crawford et al., 2000), an entirely different algorithm, or one so distributed and nonmodular so as to preclude meaningful recognition (Robinson, 1992; Stein, 1992), Most investigations of artificial neural networks have emphasized the latter possibility. Indeed, except where specific order was imposed on the network by the investigators themselves (e.g., Robinson, 1992), it has been asserted that neural nets do not employ coordinate systems or any other recognizable form of modular representation (Robinson, 1992).

However, this is at odds with the actual functional neural anatomy of, for example, the brainstem saccade generator, which parcels its tasks into relatively neat functional modules (Crawford and Vilis, 1992; Quaia et al., 1999). What is the source of this discrepancy? Does physiological modularity arise from developmental rules that are lacking in standard network models, or could the training of artificial neural networks also result in forms of functional modularity that are critical to their workings' and yet masked to casual inspection? These are crucial issues for anyone interested in using neural networks to understand brain function, the mechanisms of representation in distributed neural nets, and their relationship to the nature-nurture debate.

The aim of the current investigation was to train an otherwise unconstrained neural network model to perform the transformation from visual displacement signals to motor saccade commands and then analyze the network to see how this was accomplished. Since the correct transformation can be modeled explicitly with the use of known representations and algorithms (Crawford and Guitton, 1997), a successfully trained network must somehow be accomplishing the same thing, albeit not necessarily in the same way. Based on the recent tradition in modeling neural nets (Robinson, 1992; Stein, 1992) we expected to observe homogenously distributed representations in our network simulations but were surprised to find instead that our network achieved a level of functional modularity reminiscent of that observed at the tip of the physiologist's electrode. More important, these modules interacted through a specific algorithmic mechanism with intriguing implications for the general mechanisms of visuomotor transformation.

2. Methods

2.1. Theoretical Background

Oculomotor physiologists use the term *retinal error* to mean the angular displacement between current gaze direction and desired gaze direction, as signified by the retinal site stimulated by light from the desired target. Motor error similarly signifies the motor command sent to drive the actual movement. Retinal error has been implicitly assumed to be geometrically synonymous with motor error and has often been modeled as such (Jürgens et al., 1981; Waitzman et al., 1991; Raphan, 1998; Moschovakis and Highstein, 1994). Recently, however, it has been pointed out that when the complete 3-D geometry of the eye and saccades are considered, motor error must be considered a geometrically distinct quantity from retinal error (Hepp et al., 1993; Crawford and Guitton, 1997; Klier and Crawford, 1998).

For example, if the eye is at a special reference position called *primary position* (described below), a horizontal target will evoke a purely horizontal retinal error, and an eye movement based on the corresponding horizontal motor error will acquire that target. However, when the eye moves to a nonhorizontal position, a purely horizontal retinal error will require a nonhorizontal motor error to generate an accurate saccade that obeys Listing's law. Failure to account for this relationship would result in position-dependent errors in saccade direction (Crawford and Guitton, 1997), which are not observed in actual saccades (Klier and Crawford, 1998).

To understand why this is the case, it is helpful to review Listing's law. Primary position is a unique 3-D eye orientation often used as a reference from which to describe the relative eye-rotation vectors to other positions. Listing's law states that these rotation vectors will all lie in a head-fixed plane (Listing's plane) that is orthogonal to the line of sight at primary position (see Tweed and Vilis, 1990a). Whereas the retinal error vector is 2-D and defined in reference to the eye, the motor error vector is 3-D and defined in reference to the head since the eyes move with three degrees of freedom relative to the head and since motor error vectors normally specify saccades that lie within the head-fixed Listing's plane. Thus, although motor error is dependent on retinal error, they are not identical (Crawford and Guitton, 1997; Klier and Crawford, 1998). Therefore, to execute a saccade, retinal error in an eye-centric frame must first be converted into motor error in a head-centric frame. Moreover, it has been demonstrated that this transformation is essentially independent of plant characteristics—in particular the existence of fibro-muscular "pulleys" (Demer et al., 1995; Quaia and Optican, 1998) because they cannot solve this problem without simultaneously disrupting Listing's law (Crawford and Guitton, 1997).

Crawford and Guitton (1997) modeled this reference frame transformation by first converting the angular values for horizontal and vertical retinal error into target direction in eye coordinates. They then rotated this target direction vector by the inverse of 3-D eye position to produce desired gaze in head coordinates. The desired gaze vector was then input to a Listing's law operator (Tweed and Vilis, 1990a) that outputs desired eye position in Listing's plane. At this point, desired eye position was subtracted from initial eye position to derive the required change in eye position, or motor error. The first aim of the current study was to train a neural network to perform the same overall transformation.

2.2. Model

Whenever one models the brain, it is important to define the correspondence between the model and the system being modeled, including the inevitable limitations of the model. Although the neural mechanism of the saccadic reference frame transformation have not yet been identified, we postulate that it must be complete before activation of the short-lead burst neurons in the brainstem reticular formation because the latter appears to utilize a head-fixed 3-D coordinate system (Henn et al., 1989; Crawford and Vilis, 1992; Scherberger et al., 1998). Several recent studies have suggested that these neurons encode the derivative of 3-D eye orientation (Crawford and Guitton, 1997; Quaia and Optican, 1998; Hepp et al., 1999). We therefore trained our network to output the signal that would be appropriate to drive such neurons—that is, a 3-D vectorial change in eye orientation, defined within head-centric coordinates aligned with Listing's plane (Crawford and Guitton, 1997).

The input signal to our neural net was target direction in eye coordinates, more commonly known as *retinal error* (Crawford and Guitton, 1997). It is currently unknown how far downstream such codes might

persist within the actual visuomotor transformations of the brain. However, for the time being we will take the working hypothesis that it may be coded as low as the output signal of the superior colliculus (this question is addressed more thoroughly in the discussion section). If so, then this retinal error may be coded topographically, requiring a *spatiotemporal transformation* downstream (Tweed and Vilis, 1990b). In this case, it could be very difficult to separate this mechanism from the 2-D to 3-D and the reference frame transformations with which we were concerned. Moreover, one of the goals of this study was to avoid biologically based input constraints that might trivially lead to physiologically realistic organization in the hidden-unit solution set (Robinson, 1992). Therefore, as our visual input signal, we simply used two signals representing the already segregated components of 2-D retinal error. In this respect, this study differs from previous studies that looked at position modulations on neurons with sensory-receptive fields in a retinotopic map (Zipser and Andersen, 1988;Van Opstal and Hepp, 1995; Krommenhoek and Wiegerinck, 1998).

Another simplification is that we allowed units to "fire" positively and negatively to represent the push-pull organization seen in the brainstem oculomotor system (Robinson, 1981). Finally, since it is a trivial matter to use the motor error signal that is output from this model to drive our saccade burst-generator model and eye plant (Crawford and Guitton, 1997), we left the latter transformations out of the results described below. However, for investigators in the oculomotor field it is worth bearing in mind that the complete version of this model includes a displacement feedback loop (Jürgens et al., 1981) whose *medium-lead burst neuron* output encodes the derivative of 3-D eye orientation. The latter signal can then be used to drive a pulleylike plant that implements the saccade axes tilts required for Listing's law (Crawford and Guitton, 1997; Quaia and Optican, 1998) or can be modulated by eye position signals to neurally implement these tilts (Crawford, 1994; Crawford and Guitton, 1997).

2.3. Network Architecture

Figure 1A represents the three-layer learning neural network used in this investigation. The open circles (○) represent units that are connected by lines of varying thickness, representing the actual connection weights in one particular network. The top three units of the input layer encoded torsional (T), vertical (V), and

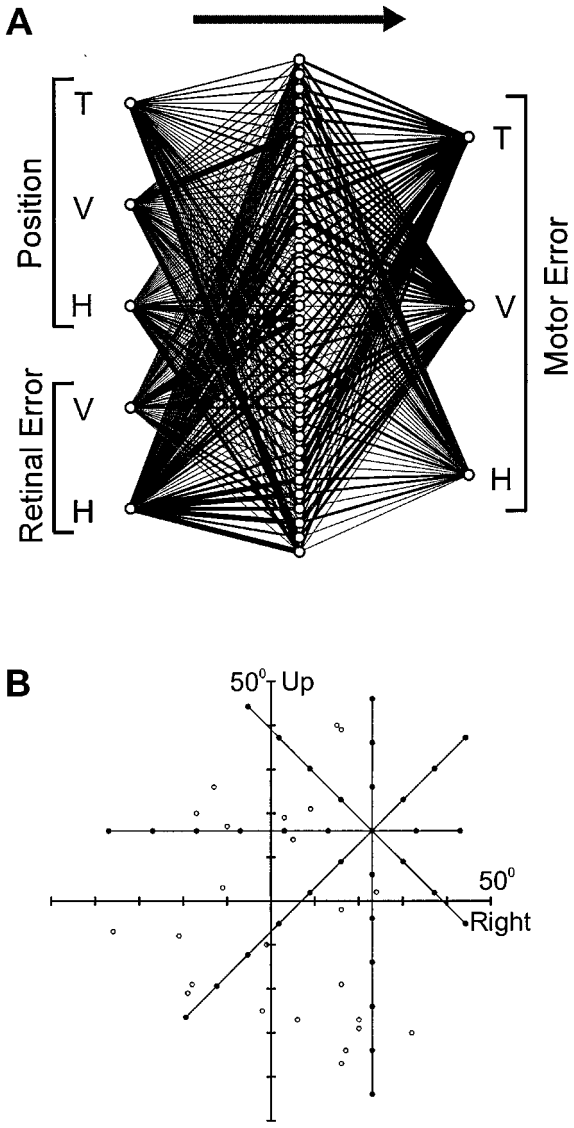

Figure 1. Neural network configuration and training set with exemplar retinal error pattern. **A:** The typical three layer fully connected neural network (N1), where ○ indicates an artificial neuron and line thickness indicates relative strength of connection. Input layer (left-hand layer) with (T)orsion, (V)ertical, and (H)orizontal components of eye position (top three values) and V and H components of retinal error (retinal error). Hidden layer (middle layer) is fully connected to input and output layers. Output layer (right-hand layer) with T, V, and H components of motor error (motor error). **B:** Standard Cartesian graph with 50 degrees maximum oculomotor range (OMR). Increments on graph are 10 degrees. ○ indicates randomly selected initial eye positions. (—•—) indicates the up to six retinal errors along each of the indicated directions associated with each of the initial positions. Each retinal error along each direction is 10 degrees greater than the former to a maximum of six retinal errors *where possible* to remain within the 50 degree OMR.

horizontal (H) components of initial eye position, while the bottom two units encoded the vertical and horizontal components of retinal error. The middle layer usually contained 35 units (see next section), each of which connected with every input unit and every output unit. The output units represented the torsional, vertical, and horizontal components of 3-D motor error.

Information flowed through the network in only one direction, from the input through the hidden to the output layer. To avoid introducing systematic bias in the networks' beginning state, each connection weight was randomly intialized (to a value within $\pm.1$). Also, the output of each neuron was limited with the use of a sigmoid shaped function: a standard practice that allows neural networks to approximate most functions. To support training in the networks we used the standard back-propagation algorithm. We also used an incremental weight update procedure. This means that the error was calculated and weights adjusted after the presentation of each training vector. We also used a momentum term with the back-propagation training algorithm. This modification changes the weight-adjustment procedure by adding a fraction (in these networks 10%) of the previous weight adjustment to the current adjustment as calculated by the training algorithm. This reduces the incidence of the network becoming stuck in local minima—that is, suboptimal solutions (Freeman and Skapura, 1991; Rumelhart et al., 1986). (See the appendix for formulas relating to the network.)

We chose back-propagation as a training rule because it fits with the goals of this study, which were not primarily to find the global, optimal solution—for which rules such as the genetic algorithm may be superior—but rather to describe the effects of training through sensory feedback and successive approximation on the functional organization of the network.

2.4. Network Training

The network was trained on a group of input patterns specifying initial 3-D eye position in craniotopic coordinates and 2-D retinal error vectors in eye coordinates. Figure 1B illustrates the eye positions used and a typical set of associated retinal error vectors. The input patterns consisted of 25 initial eye positions (○) randomly presented within a 40 degree oculomotor range. The torsional component (not shown) was always set to zero—that is, all initial eye positions were within Listing's plane. The origin in Fig. 1B represents primary position. Associated with each initial eye position were up to 48 retinal error vectors. These retinal error vectors were of six magnitudes (10, 20, 30, 40, 50, and 60 degrees) along each of eight directions *where possible* without exceeding ± 50 degrees from the origin. For example, the upper right-hand quadrant of Fig. 1B illustrates a chosen position with the associated retinal errors (●) connected by lines radiating in the eight indicated directions. Each of the retinal errors along any one of the directions was presented sequentially for each initial eye position (○), from 10 to 60 degrees (where possible without exceeding ± 50 degrees from the origin) although each direction was presented randomly.

A corresponding group of correct desired target vectors specifying 3-D motor error vectors in head-centric coordinates completed the training set. These ideal motor error vectors were computed for each input set by converting the angular components of retinal error into an eye-centric pointing vector and position into a quaternion and then inputting these values into the algorithm described above and in more detail in Crawford and Guitton (1997). The output was then converted into a 3-D displacement vector with angular components.

The network was then trained to minimize the error between its actual output and these correct motor error vectors, by a supervised training process resembling the actual calibration process whereby saccades learn to be accurate through visual feedback (Fuchs et al., 1985; Optican and Miles, 1985). Thus, the network was trained to map a variable set of retinal errors and positions onto the correct pattern of motor errors for accurate saccades that obey Listing's law.

Using the training set outlined above, we trained several networks in which we varied the number of hidden units. We began with a minimum of five hidden units and increased that number in increments of five up to a maximum of 40 hidden units. Network performance was assessed by monitoring the squared-error between the desired output and the actual output of the network. Figure 2 shows a typical output error curve for networks with the various number of hidden units. As shown in the figure, with five units in the hidden layer, average output error reached a stable minimum of approximately 21 degrees within 200 epochs (an epoch represents one presentation of the entire training set. With 10 hidden units, output error was reduced to approximately 8 degrees. With an increase in the number of hidden units to 35, output error

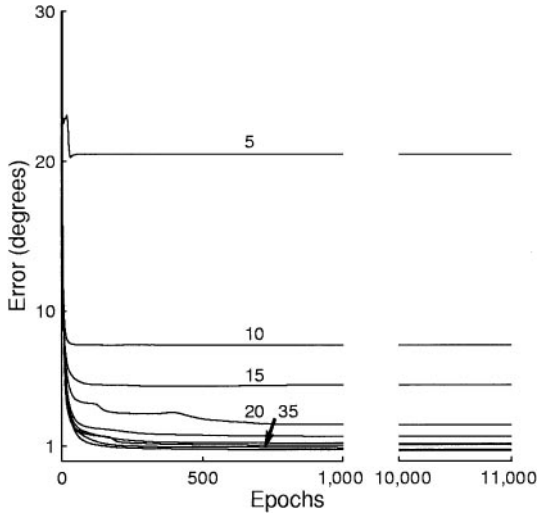

Figure 2. Output error (difference between actual and ideal performance) from representative networks with varying numbers of hidden units. Numbered traces indicate quantity of hidden units used in a representative network. Hidden units in these networks varied from 5 to 40 in increments of 5. Note that only incremental reduction of error was achieved by networks with greater than 20 hidden units and that a stable minimum output error for all networks was reached within approximately 1,000 epochs.

was further reduced to within 1 degree of ideal performance. Since a network with 40 hidden units produced only an incremental improvement over one with 35 and since a 1 degree error was within the range of experimentally observed errors (Klier and Crawford, 1998), we used only those networks with 35 hidden units (a total of 13 networks (N1–N13)) for further testing and analysis. Figure 2 also shows that a stable minimum error for these networks was achieved within approximately 1,000 epochs. Nevertheless, training for all networks was allowed to continue for another 10,000 epochs to ensure that a stable minimum for each network had actually been reached.

2.5. Network Testing

After training was completed, the networks were tested with several sets of corresponding input and target patterns. The first set consisted of the patterns used for training. Testing with this set of patterns allowed us to assess how well the networks learned the task. Another set of test patterns, that were never used during training, consisted of a randomly generated set of input vectors (and associated target vectors) and was used to assess the generalizeability of the networks' solutions. In addition to these, we also used a set of testing patterns with specified positions along the diagonal from the upper left, through the origin, to the lower-right of a Cartesian coordinate system (see the coordinate system in Fig. 1B for an example) and a set of four eye positions originating at -30, -15, 15, and 30 degrees along both the x and y axes. Testing with these patterns, which were never used for training, allowed us to assess a network's performance along tertiary and secondary eye positions, respectively.

We also tested the network with *stimulation* and *lesion* trials. Stimulation of a site upstream from the network encoding retinal error was modeled by varying eye-position input to the network of any one trial while holding retinal error constant as if electrically microstimulating an upstream visual site coding a constant retinal error. For example, if a particular unit had a large weighting for the vertical component of retinal error and a zero weighting for the horizontal component of retinal error, its retinal sensitivity would be represented by a vertical vector. The lesion trials were conducted by setting the outputs of selected hidden units to zero (as described in more detail in the Results section).

Finally, individual hidden unit characteristics were also examined by plotting the direction and magnitude of their *sensitivity vectors*. We defined a hidden unit's sensitivity vector as the direction and magnitude of the weights associated with a particular hidden neuron. Thus, a hidden unit had three sensitivity vectors: two input vectors (components of position and retinal error) and one output vector (components of motor error). The results of network training, testing, stimulation, and lesioning were assessed by quantifying the differences between the actual output of the network and ideal motor error as a function of direction and magnitude.

3. Results

This section has two purposes. First, to illustrate the position dependent relation between retinal error and motor error in accurate saccades. Note that this has already been confirmed experimentally (Klier and Crawford, 1998). The second, and more important purpose was to describe how well our network learned this geometry.

Figure 3 shows the output of network N1 from nine selected positions that *were not used in training*. In this and other plots, the horizontal and vertical components of 3-D angular position are plotted in Listing's coordinates, but we dispensed with the right-hand convention

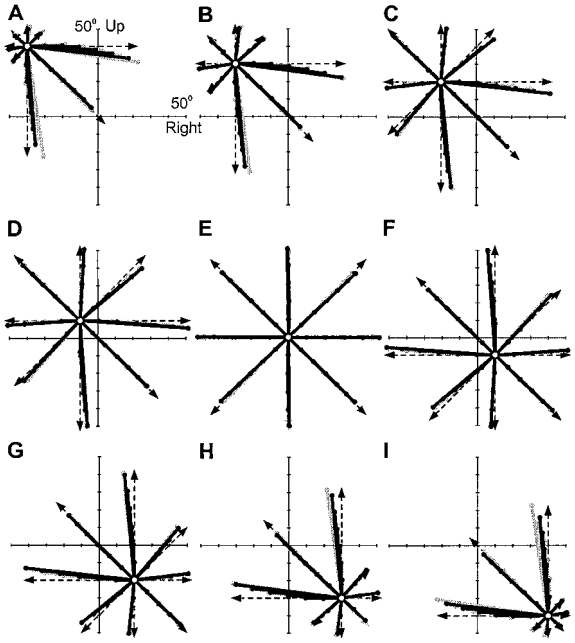

Figure 3. Nine selected positions used to test network performance. All initial positions (○) are within 50 degrees OMR shown on standard Cartesian graph (increments on graph are 10 degrees). Retinal errors indicated by dashed lines ending in arrows are provided for reference. Actual network output: solid dots connected by thick dark lines. Ideal network output shown in thick light lines. Initial positions are in 10 degree increments along the diagonal starting at **A** (-40°, 40°), to **E** (0°, 0°), to **I** (40°, -40°). Initial position on the origin (**E**) corresponds to *primary position* (see text). Note that actual network output and ideal output tilt away from retinal error depending on initial position and magnitude of retinal error. Actual network output closely follows ideal and often obscures it.

for simplicity (that is, up is up, and right is right). The nine positions are staggered in 10 degree horizontal and vertical steps along the diagonal from the upper left (40°, 40°) in Fig. 3A to the lower right (-40°, -40°) in Fig. 3I. Input to the network for these positions was 2-D retinal error in an *eye-centric* frame (⟶). Note that these schematically plotted retinal error vectors do not indicate actual target direction in the head-centric coordinate system of the figure, but they do show the direction that the eye would move if retinal error was not compensated for by eye position (Crawford and Guitton, 1997). Each initial position in Fig. 3 is associated with six retinal errors at 10 degree intervals of magnitude (where possible without exceeding ± 50 degree oculomotor range) for each of our eight standard radial directions and the actual motor error outputs of the network (—●). For comparison, the computed ideal motor error showing the correct saccade directions is indicated by light lines, although these are often obscured by the overlying actual motor error output by the network.

Motor error diverged from the supplied retinal error as a function of position eccentricity, direction of retinal error, and magnitude of retinal error. For example, in Fig. 3A, motor error tends to converge toward the origin. Saccades at the opposite corner (Fig. 3I) showed the symmetrically opposite effect, whereas saccades at intermediate positions (Fig. 3B–H) showed similar patterns but with intermediate magnitudes. Only in Fig. 3E, where the initial position was at the origin (primary position), was there no divergence of motor error from retinal error as expected. Note that in each of the panels of Fig. 3, the actual output of the network and the ideal output are almost indistinguishable.

In summary, motor error and retinal error aligned with each other at primary position (Fig. 3E) but showed increasing divergence as the initial position grew more eccentric. These illustrative initial positions show that the networks were clearly able to generalize this position-dependent pattern from the initial training set to a set of new positions which were never used during training.

In Fig. 3, all of the eye positions except primary position (Fig. 3E) are *tertiary*—that is, oblique. Such eye positions involve the so-called *false torsion* first described by von Helmholtz (1925). For example, if a movement induced a counterclockwise false torsion of the eye (as would occur at the position illustrated in Fig. 3A), then *space horizontal* would be rotated clockwise with respect to the horizontal retinal meridian. Thus, a purely horizontal retinal error may reasonably be expected to induce a nonhorizontal motor error at such positions. However, false torsion explains only part of the position-dependent deviation from retinal error seen in these data (Crawford et al., 2000). To demonstrate this, we also simulated saccades from secondary (that is, vertical and horizontal) positions in Listing's plane where there is no false torsion (Fig. 4). Note that, again, these positions were not used during training and thus were never seen by the network prior to this test.

Figure 4 shows the neural network output (dark lines) with initial eye positions arranged along the vertical axis in Fig. 4A and the horizontal axis in Fig. 4B. At these initial positions, retinal errors from -50 to 50 degrees (at 10 degree intervals) were input along the orthogonal axes (horizontal in A and vertical in B).

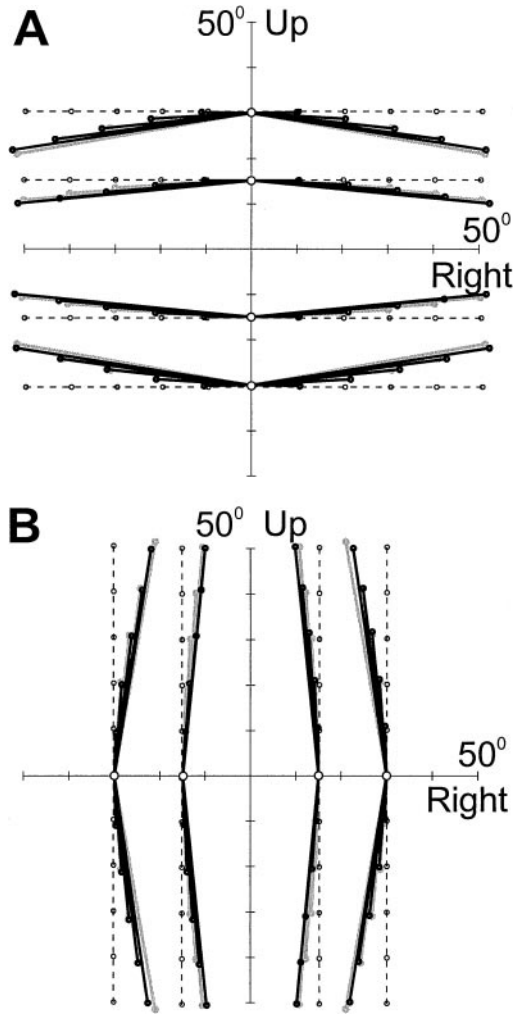

Figure 4. Positions on primary axes (horizontal and vertical) used to test network output. **A:** Initial positions arrayed along y-axis at 15 degree intervals from −30 to 30 degrees on 50 degrees OMR Cartesian graph. Associated with each position are 10 purely horizontal retinal errors (- - -○- - -) five from −50 to 10 degrees and five from 10 to 50 degrees at 10 degree intervals. Actual network output indicated by dark lines: (—●—). Ideal network output indicated by light lines (—●—). **B:** Same conventions as in **A** but initial positions are arrayed along the *x* axis. All positions in the figure are *secondary* and cannot induce false torsion. Therefore, false torsion cannot be the cause of the pattern of deviation from retinal error seen in the ideal and actual output of motor error. Note that degree of deviation increases with eccentricity from primary position (origin) and magnitude of retinal error.

Despite the absence of false torsion, a deviation between retinal error and motor error still occurred depending on the magnitude of retinal error *orthogonal* to eye position. For example, with eye position at 30 degrees *vertical*, a 50 degree *horizontal* retinal error (dotted lines) produced a network output that deviated vertically by about 10 degrees. Note that this deviation was always in the direction *opposite* to the direction of initial position. A similar but reduced pattern was seen at an initial position 15 degrees upward. At each position, increasing the magnitude of retinal error increased the deviation in the network output. The same patterns of deviation were also seen (but in the opposite direction) when eye position was initialized at downward positions. Moreover, a similar pattern of motor-error deviations was seen when eye position was initialized in the horizontal direction with vertical retinal errors (Fig. 4B). The main point to note is that whenever the eye moves from primary position, retinal error cannot be trivially mapped onto motor error without a concomitant reduction in saccade accuracy. Again, the network learned to correct for this.

Next, these observations were quantified across networks by testing each network with a set of five initial positions that were not used during training: (1) primary position (0°, 0°), (2) up-right (20°, 20°), (3) up-left (−20°, 20°), (4) down-left (−20°, −20°), and (5) down-right (20°, −20°), Fig. 5B–F. Each of these positions was associated with eight retinal error vectors of 30 degrees magnitude oriented in the same radial pattern as before. The results of this test were quantified as the *mean angular difference* between the direction of actual motor error output by the networks and the direction of ideal motor error (—♦—) Fig. 5A. For reference we also quantified the angular difference between retinal error and ideal motor error (----◇----) as a measure of the error expected with zero compensation for position (Klier and Crawford, 1998). The mean error across networks for positions 1 through 5 are plotted in Fig. 5B through F respectively as a function of retinal error direction in polar coordinates. Thus, Fig. 5B–F plots actual directional error (—♦—) compared to directional error without position compensation (----◇----), across eight directions for each such position.

Figure 5B shows position 1 (0°, 0°). At primary position there is no need to compensate for position effects, and this fact is reflected in the pattern of error, which remains at zero for ideal motor error and close to zero for actual motor error. In the other plots of Fig. 5, the pattern of error for a lack of position compensation (◇) oscillates between ±6 degrees for 30 degree saccades. For example, in Fig. 5C, for a 0 degree angle of retinal error (purely horizontal and rightward) the error starts

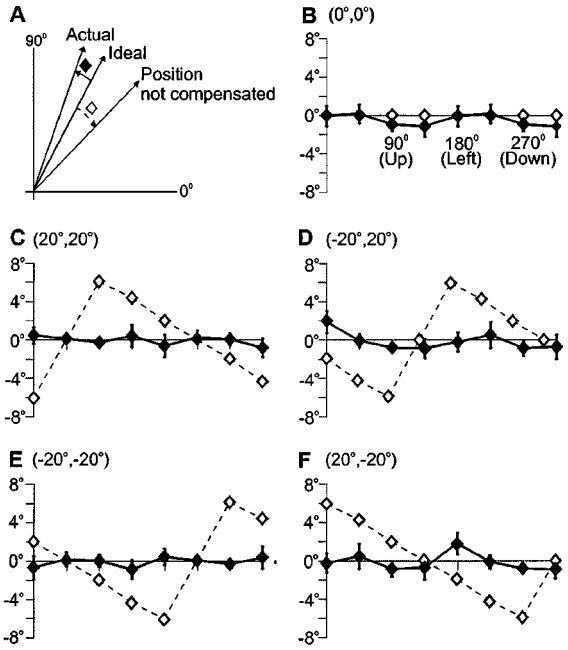

Figure 5. Angular difference measures across networks. **A:** Schematic showing geometric relationship of the angular difference measures plotted in **B–F**. The vector labeled "Position not compensated" is the supplied retinal-error vector. The vector labeled "Ideal" is a hypothetical ideal motor-error output vector, and the vector labeled "Actual" is motor-error output from a hypothetical neural network. The mean angular difference measure between retinal error and ideal motor error is indicated by (----◇----) in **A** through **F**. The mean angular difference measure across networks between actual motor error produced by a network and ideal motor error is indicated by (—♦—) in **A** through **F**. Angular difference measures were taken from five initial positions on a Cartesian graph: **B** (0°, 0°—primary position), **C** (20°, 20°), **D** (−20°, 20°), **E** (−20°, −20°), **F** (20°, −20°). Each position has associated with it eight retinal errors of 30 degree magnitude in the same radial pattern as illustrated in Fig. 1B. The *x* axis in the plots **B–F** is labeled in accordance with the angle between the axis and the retinal error with 0 degrees indicating a purely horizontal and rightward retinal error. The angle between the axis and the vector increases in a counterclockwise direction. Thus, 90 degrees indicates a purely vertical and upward retinal error. Note that perfect agreement between the angular difference measures would cause the values to lie along the *x* axis. For example, with perfect agreement between ideal motor error and retinal error (----◇----), the trace would lie along the *x* axis as in **B** (as expected in primary position (0°, 0°)).

out at −6 degrees, progresses to 6 degrees for a 90 degree angle of retinal error, and then decreases again to −2 degrees for a purely downward angle of retinal error. Note, however, that the actual output of the networks (♦) does not follow this pattern. Rather, the actual output remains relatively flat, indicating a near-zero error between the correct ideal behavior and the actual behavior across networks. A similar pattern can be seen in each of the other panels. Thus, all of the networks learned the correct transformation to motor error and were able to generalize this learning to all of these positions and all directions of retinal error. In addition, the torsional behavior (not shown) of these networks was also highly precise, showing deviations from Listing's plane of only −.229 degrees, ±.003 (mean ±SD) across all networks.

3.1. Stimulation Trials

Next, we performed the equivalent of stimulation of a site on a retinotopic map encoding retinal error upstream from our network. We did this by testing the network using an input set (never used during training) where retinal error remained constant and only the initial eye position varied. Initial eye positions were randomly chosen to remain within a ±50 degree oculomotor range. Figure 6 displays the results of these trials where the ideal response is shown in light lines, while

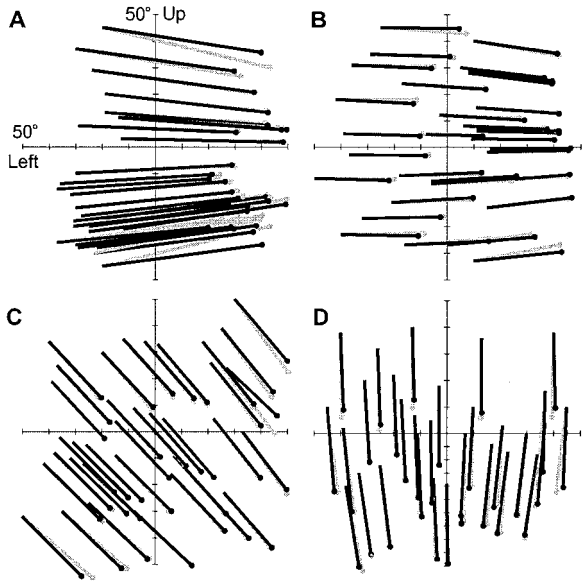

Figure 6. Stimulation trials network 1 (N1). Dark lines: Actual motor error output from network. Light lines: Ideal motor error output. **A:** 60 degrees purely rightward retinal error (not shown). Note converging pattern of actual *and* ideal motor error. Network output follows ideal output fairly closely. **B:** 30 degrees purely rightward retinal error. **C:** 30 degrees downward and rightward retinal error. **D:** 30 degrees purely downward retinal error. Note converging pattern of motor error is dependent on initial position and magnitude of retinal error (see text).

the actual response of the network to the stimulation of a constant retinal error is shown in dark lines. Figure 6A shows a 60 degree rightward stimulation. Note that as the initial position becomes more distant from the *x* axis the ideal response tilts at a greater angle towards it. The tilting pattern of the ideal response (light lines) is closely matched by the actual output of the network (dark lines).

Figure 6B shows a similar pattern of response for ideal and actual output from the network for a smaller 30 degree horizontal retinal error. Note that the tilting effect is less than that seen in Fig. 6A. This lessening of the tilting effect is due to the smaller retinal error used in this data set, as expected from the pattern described above. Figure 6C and D demonstrates that this pattern of convergence toward primary position holds for any retinal-error direction. Looked on as a complete pattern, these diagrams show a converging effect of motor error that is based on both initial eye position and retinal error magnitude and that bears a strong resemblance to studies involving actual stimulation of the superior colliculus and frontal cortex (e.g., Freedman and Sparks, 1997; Schlag and Schlag-Rey, 1987; Russo and Bruce, 1993).

To quantify this convergence, we used the *orbital perturbation index* developed by Russo and Bruce (1993) to quantify the eye-position-dependence of saccades evoked by stimulation of the frontal cortex. This index was computed separately for the horizontal and vertical components of eye movements, as the linear regression of saccadic displacement from each initial eye position (K_h and K_v, respectively). An index value of 0.0 would indicate a constant vector, while a score of -1.0 would indicate goal-directed saccades that end at the same position regardless of the initial eye position. This analysis was performed for simulated stimulations of 30 and 60 degree retinal-error sites across all 13 networks, quantified for both ideal and actual saccades. The mean results and standard deviations for actual network output and ideal output across networks were for 30 degree stimulation: actual $= -0.073$ (SD $= 0.005$), ideal $= -0.084$ (SD $= 0.000$). For 60 degree stimulation the results were actual $= -0.245$ (SD $= 0.008$), ideal $= -0.306$ (SD $= 0.000$).

3.2. Hidden-Layer Analysis

To assess how the network may be accomplishing the above transformations, we looked at the input and output weightings of units in the hidden layer of the trained network. Recall that each input that a hidden unit received was multiplied by a scalar value roughly corresponding to a synaptic strength, whereas the weights of its connections to the output units defined its motor tuning. We initially hypothesized that the network accomplished its task through certain three-way correlations between weight components of position, retinal error, and motor error. However, an analysis of correlations between all combinations of position, retinal error, and motor error weightings across all of the hidden units in each network revealed little about how the network accomplished the task. Some correlations were slightly higher than others, but all were weak. Mean correlations between components ran from $-.03$ to .07 with none of them being statistically significant. This seemed to suggest that the solution was completely distributed so as to preclude any meaningful analysis (Robinson, 1992). To investigate further, we next tried a more physiological approach: the equivalent of *cell recording*. That is, we analyzed each hidden unit by studying its sensitivity vectors.

By a *sensitivity vector* we mean the vector defined by the components of a units' weights. For example, each neuron in the hidden layer received five weighted inputs signals from the input layer: three for position and two for retinal error. We treated these signals as the components of two vectors: a 3-D position vector (Fig. 7A: blue vectors) and a 2-D retinal error vector (Fig. 7A: red vectors). Similarly, the three weighted signals to the output layer were also considered as the components of a 3-D motor error vector (Fig. 7A: black vectors). Figure 7A shows these weight vectors for each of the 35 hidden units of network N1, for which only the horizontal and vertical components are plotted; the third dimension of position and motor error will be treated in a later section.

Figure 7A shows the sensitivity vectors of each hidden unit within network N1 labeled according to the following scheme. Visual inspection of such data suggested to us that network units may have formed specific task groups. In many units (1, 3, 4, 5, 6, 7, 8, 9, 12, 13, 15, 16, 20, 21, 22, 24, 33), the retinal-error vectors and the motor-error vectors were aligned, and the magnitudes of their retinal-error components were about half the size of their motor-error components. In addition, these units tended to have a very small position component. These units seemed to be directly mapping retinal error onto a motor response regardless of position, like many of the vector-displacement models (Jürgens et al., 1981; Raphan, 1998), so we named these

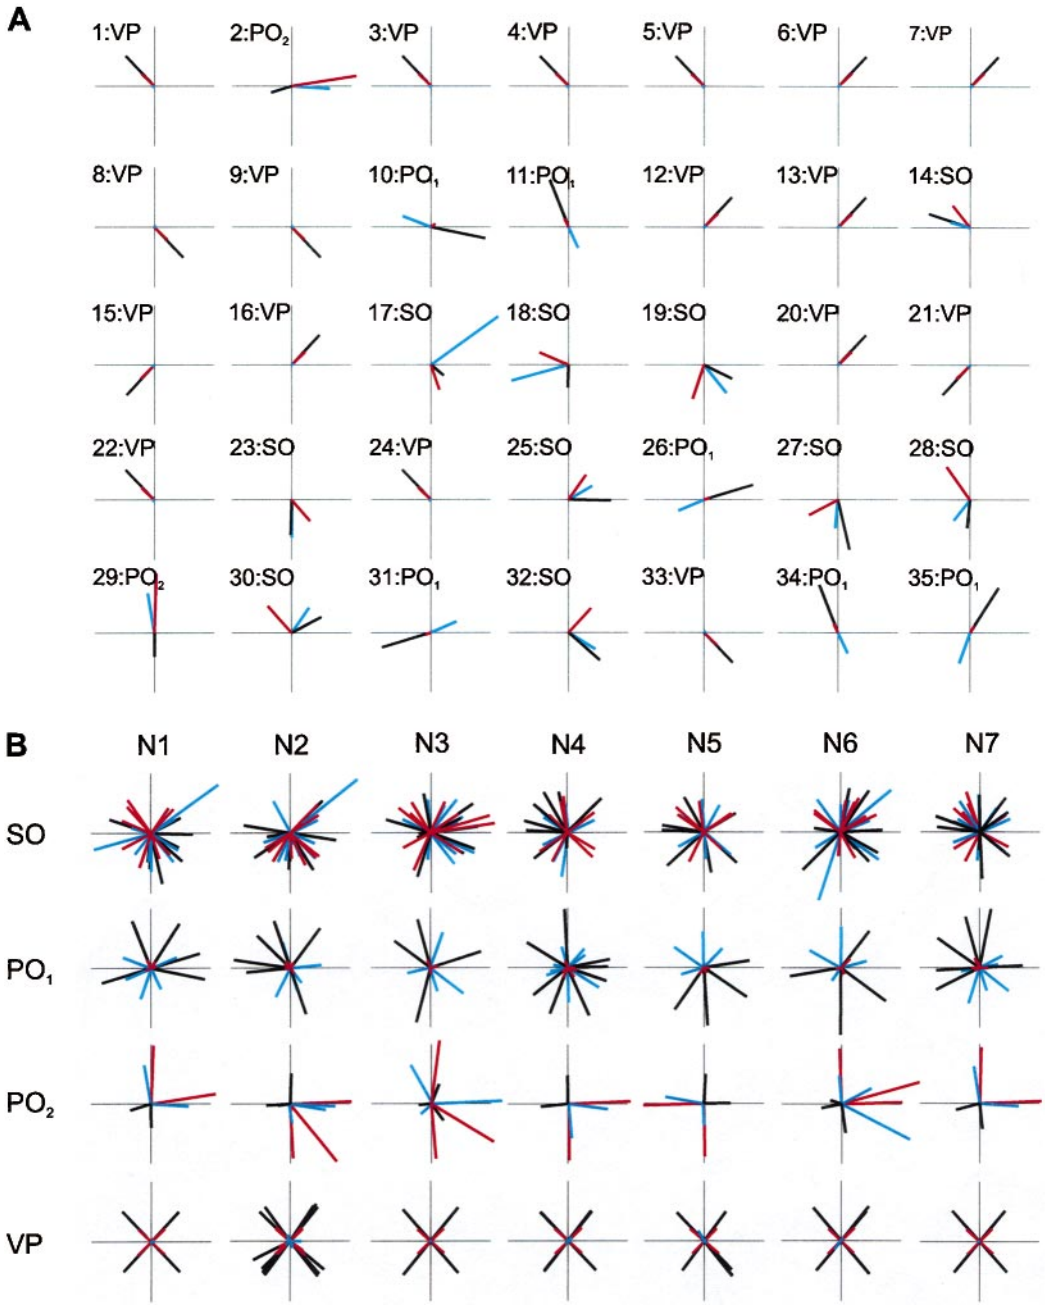

Figure 7. Geometric interpretation of hidden unit connection weightings as input/output sensitivity vectors. **A:** Network 1 (N1). Angle and magnitude of sensitivity vectors are measured from the origin of a Cartesian axis. Individual plots are labeled according to hidden unit number and group type, respectively. VP: Vector-propagation group. PO_1: Position-opposite group 1. PO_2: Position-opposite group 2. SO: Semiorthogonal group. The different sensitivity vectors are indicated by color: red is retinal error, blue is position, black is motor error. **B:** All of the above groups (same conventions) superimposed onto single plots by group, across seven networks (N1–N7). Note that the vector-propagation group forms a coordinate system rotated ~45 degrees from a standard Cartesian graph.

units the vector-propagation group. Other units (10, 11, 26, 31, 34, 35) tended to have position and motor errors in opposite directions with a small retinal-error component. A smaller subset of these units also had position tuning opposite to motor error tuning, but they had a large retinal error vector (2, 29). We hypothesized that these latter two groups provided the main position dependent effect on motor error output by the network (Figs. 4 and 5), since the position-dependent effect on motor error is in the direction opposite of the orthogonal component of position. Thus, we named these two groups position-opposite$_1$ and position-opposite$_2$ respectively, since they appeared to perform similar functions. The distinction between these two groups was to prove useful in further analysis. Finally, the remaining units (14, 17, 18, 19, 23, 25, 27, 28, 30, 32) tended to have their motor-error and retinal-error vectors in quasi-orthogonal directions, while the magnitude of their motor error, retinal error and position components tended to be more variable than the other groups. We hypothesized that these units may provide a fine tuning of the response in all available areas of the solution space, perhaps modifying the amount of position compensation depending on the length of the retinal error signal. These we named the semiorthogonal group. The labeling in Fig. 7A reflects this classification scheme applied to N1.

To systematize our initial observations of hidden-unit sensitivity vectors, we developed quantitative criteria to classify the above described units (see Table 1). Through trial and error, we found that our algorithm broke the weighting data down into our four classes and captured 98% of the units in all networks. Using these criteria, we found that the number of each different type of unit was quite consistent between networks: vector-propagation: 17.1 $\pm$1.5 (mean $\pm$SD); position-opposite$_1$: 5.9 $\pm$1.1; position-opposite$_2$: 2.7 $\pm$.6; semi-orthogonal: 9.4 $\pm$1.1. (Of course, there was no correlation between which hidden unit became which type because the units' initial weights were randomized). Moreover, when we separated the units of each network into these classes and then superimposed their sensitivity, as shown in Fig. 7B for networks N1 to N7, the group distributions of these vectors were quite consistent between networks. We will return to Fig. 7B in a later section on population coding, but first let us consider the objective validity of this classification scheme in more depth.

To visualize the separate clustering of these hidden-unit classes, we replotted their sensitivity vectors in a way that illustrates all of the data simultaneously, color coded according to the classification criteria shown in Table 1. Figure 8 shows these four classes as plotted in 3-D spaces that are defined by the angles between the eye-position, visual, and motor-sensitivity vectors (along the horizontal axes) and the magnitudes of these various vectors (along the vertical axes). The horizontal plane of the panels in the left column (A, B, C) plot the visual-motor angle versus the visual-position angle, whereas the right column (D, E, F) plots the visual-motor angle versus the position-motor angle ($\pm 180^\circ \times \pm 180^\circ$). Note that the opposite sides of each such plot are really continuous but have been separated in the process of mapping this cyclic information onto a plane. Each of these horizontal planes is then divided into a 2-D grid of 36×36 bins of ($10^\circ \times 10^\circ$). The vertical height of each bin is then specified by the summed magnitudes of all sensitivity vectors—across all networks—that fall within that bin. The magnitudes of the three sensitivity vectors have been plotted separately, so that the top row (A, D) shows the input weighting of the visual signal, the middle row (B, E) shows the input weighting of the position signal, and the bottom row (C, F) shows the motor output weighting.

For the most part, this plot resulted in clear clusters of data, with each of our classification groups forming

Table 1. Functional criteria for hidden unit classes.

Unit Type	Membership Criteria
Vector-propagation	Angle between the retinal error and motor error $<=10^\circ$ *and* the magnitude of the retinal error $>$ position
Position-opposite$_1$	Angle between the motor error and position $>=150^\circ$ *and* the magnitude of the retinal error $<$ position
Position-opposite$_2$	Angle between the retinal error and motor error $>= 10^\circ$ *and* the angle between the motor error and position $>= 150^\circ$ *and* the magnitude of the retinal error $<$ motor error
Semiorthogonal	Angle beween the motor error and position $<= 150^\circ$ *and* the angle between the retinal error and motor error $<10^\circ$

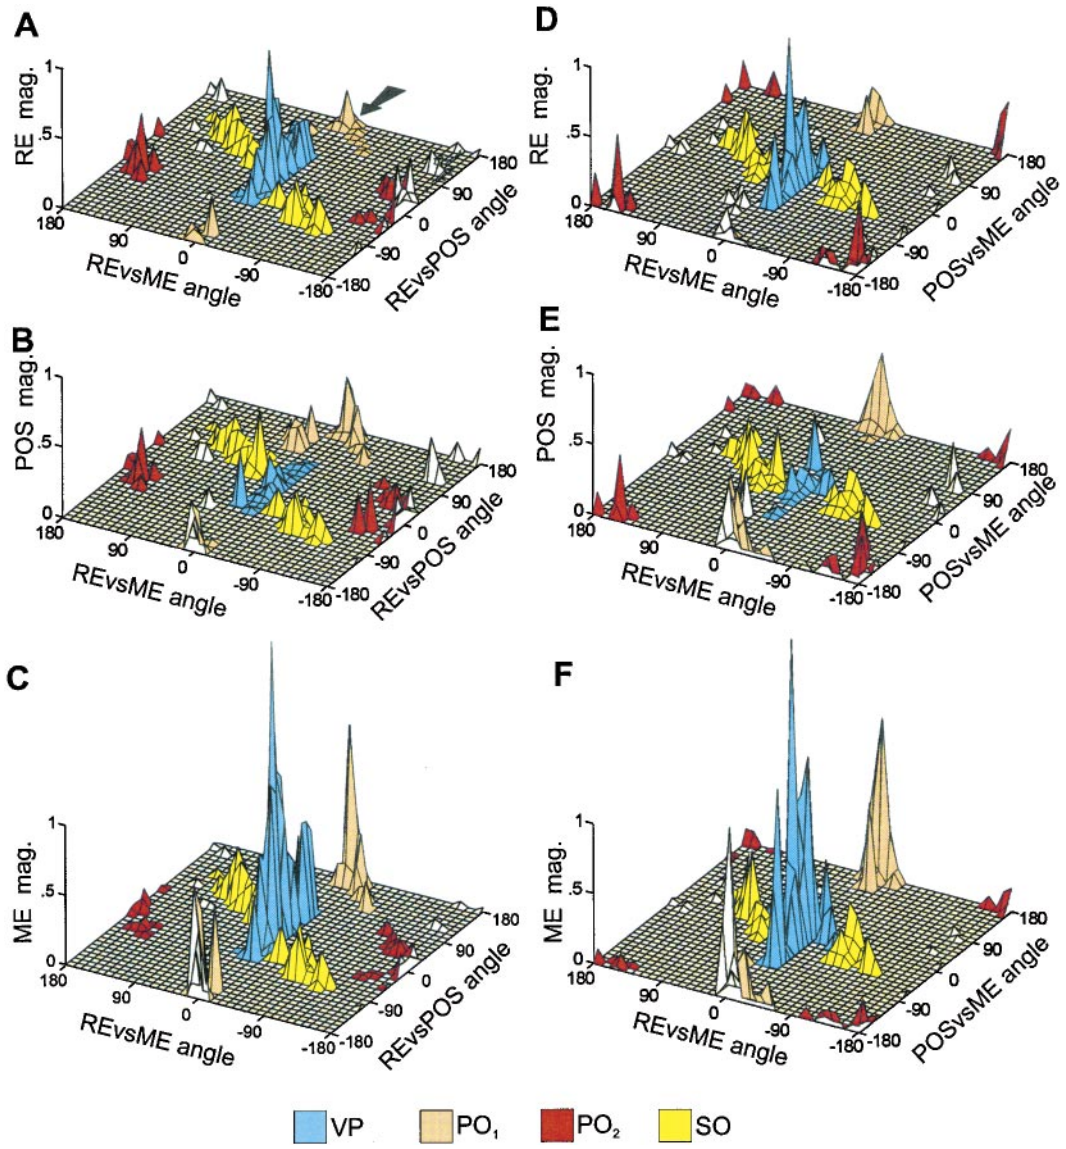

Figure 8. Three-dimensional group plots of hidden unit sensitivity vectors across networks. Left-hand column: Floor of graph formed by angles between (retinal error and motor error) and (retinal error and position), shaded by group—vector-propagation is blue, position-opposite$_1$ is orange, position-opposite$_2$ is red, semiorthogonal is yellow. Small scattered, uncolored peaks represent outlaying units that did not fall into our classification scheme and had minimal impact on behavior. *Y* axis records magnitude of sensitivity. **A:** Retinal error sensitivity. **B:** Position sensitivity. **C:** Motor error sensitivity. Arrow in **A** indicates viewing direction when comparing Fig. 8. Right-hand column: Floor of graph formed by angles between (retinal error and motor error) and (position and motor error). Other conventions the same as left-hand column.

a separate "mountain" in sensitivity vector space. The position-opposite groups 1 and 2 (orange and red, respectively) formed separate peaks near the edge of the range. At first glance each of these groups seems to be divided into two or four separate peaks at opposite edges, but note that this is just because of the wrap-around discontinuity at the edge. Accounting for this, they are really just forming one peak each. The thin central blue peaks correspond to the vector-propagation group, stretching along the line that signifies parallel visual- and motor-sensitivity vectors. Looking at any one panel, these blue peaks may sometimes be hard to distinguish from the yellow semiorthogonal range that cuts diagonally across the blue range, splaying outward

from either side. But when one inspects these ranges across rows, one can see that these two ranges behave differently—for example, the blue showing larger visual input (A, D) and less position input (B, E) than the yellow.

Since the bottom row of this plot gives the magnitudes of motor output, this row gives an idea of the relative behavioral importance of a given hidden-layer class for a given input. Thus, the huge blue peaks in the bottom row signify that for a given input, the vector-propagation units provided the greatest single source of motor output, as one might expect from a class that appears to do the main job of mapping the visual stimulus onto a spatially similar motor output. Clearly, the individual peaks corresponding to the other classes would have less weight on behavior (for a given input), as one might expect from units that are performing a more subtle, modulatory role. However. it is also important to recall that the motor outputs represented in the bottom row depend on the sensory inputs seen in the top two rows. Thus, the overall input-output relations of each cluster can only be seen by scanning up and down the columns in Fig. 8, tending to confirm our original observations based on individual units.

3.3. Cluster Analysis

Since the preceding analysis relied on visual inspection, we also performed a formal quantitative cluster analysis on the data. Using the SPSS statistical package, we performed a standard hierarchical cluster analysis using within groups linkage and the squared Euclidean distance between groups. Input to the analysis consisted of three angular values (retinal error versus motor error, retinal error versus position, and position versus motor error) as well as three magnitude measures (retinal error, position, and motor error) for all hidden units across all networks. However, since there was a discontinuity at the ± 180 degree angle, the angular values were replaced with both the sine and cosine values for this analysis.

This cluster analysis consistently arrived at four significant clusters (even for higher-order fits), and these clusters agreed substantially with our initial quantitative groupings—for example, showing 91% inclusion of the vector propagation units in cluster 1, 95% inclusion of the position-opposite$_1$ units in cluster 2, and 100% agreement of the position-opposite$_2$ class with cluster 3. However, there was one exception: the cluster analysis did not separate the vector propagation and semiorthogonal groups but instead treated them as a continuous range. Considering the subtleties noted above for distinguishing these blue and yellow ranges in Fig. 8, this is perhaps not surprising.

One possible explanation for the minor discrepancy between our initial scheme and the formal cluster analysis is that the cluster-analysis algorithm was not subtle enough to account for the second-order effects that we felt were important for distinguishing between the vector-propagation and semiorthogonal groups or, in other words, that it was fooled by the proximity of these groups in the horizontal plane illustrated in Fig. 8. Presumably, one could alter the form and weightings of the inputs to the cluster analysis until it agreed with our intuitions, but this would seem to be an exercise in self-affirmation. Therefore, to test our own intuitions more rigorously, we graphically analyzed the hidden units across all 13 networks ($13 \times 35 = 455$ units) in a way similar to that shown in Fig. 8., but this time looking at the individual units rather than entire populations.

Figure 9 illustrates the input magnitudes as a function of the angle between their sensitivity vectors—that is, retinal error versus motor error (left column) and position versus motor error (right column). The rows further quantify the properties of each unit by showing the *magnitude* of its visual-sensitivity vector (row 1) and its position-input vector (row 2). Note that these plots show individual units, categorized as vector-propagation (●), position-opposite$_1$ (◇), position-opposite$_2$ (♦), and semi-orthogonal (○) groups according to our original criteria (Table 1).

In panels A and C, both the semiorthogonal (○) and vector-propagation (●) groups are intermeshed and *seemingly* range over a continuum, so again, it is easy to see how a cluster analysis influenced by these views would put them together. However, that these are two different groups becomes apparent in plots B and D. The response magnitude of the semiorthogonal (○) group in these plots remains relatively constant and clustered near the center of the graph, while the vector-propagation group (●) forms a tight line, which is particularly separate in plot D. This suggested to us that our initial classification scheme was essentially correct, despite the one discrepancy from the formal cluster analysis. Moreover, as we shall see in the following sections, the distinction between the vector-propagation and semi-orthogonal groups proved to be important in further analysis where these groups showed very different population characteristics and functional roles.

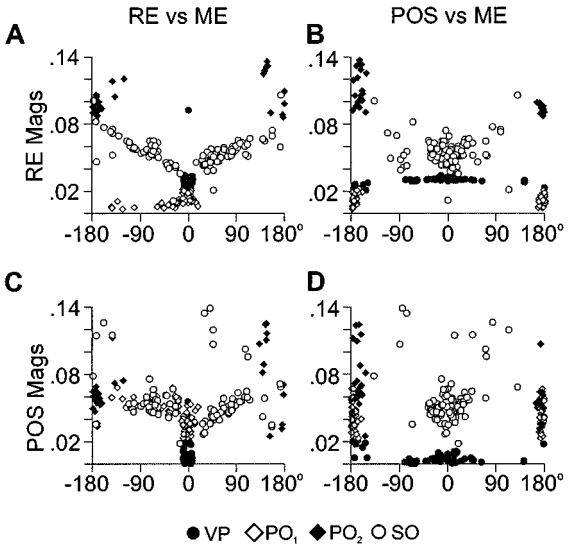

Figure 9. Sensitivity magnitudes plotted as a function of angular differences between hidden unit sensitivity vectors across networks. The *x* axis of each plot shows the angle between hidden unit sensitivity vectors: retinal error (RE) and motor error (ME)—left-hand column, and position (POS) and motor error—right-hand column. The *y* axis of each plot shows the magnitude of the input-sensitivity vectors: retinal error (first row) and position (second row). Each point represents a single hidden unit so that all hidden units for all networks are represented in each graph. Points are labeled according to our quantitative criteria: VP: vector-propagation ●. PO_1: position-opposite$_1$ ◇. PO_2: position-opposite$_2$♦. SO: semiorthogonal ○. By scanning the different graphs, it can be seen that points form clusters rather than an amorphous distribution and that groups that seem to be continuous in one view can be distinguished from each other in another view (e.g., PO_1 and PO_2 in plots B and D).

3.4. Population Coding and Coordinate Systems

Having established that the network's hidden layer could be legitimately divided into four parallel groups, we then asked how these populations coded for saccade direction. To do this, we overlapped the sensitivity vectors for each network separately for each type of group. Figure 7B shows these groupings for networks N1 to N7. Other networks (not shown) displayed very similar patterns of sensitivity groupings. The top row depicts the semiorthogonal group, the second and third rows show the position-opposite$_1$ and position-opposite$_2$ groups, and the last row displays the vector-propagation group. Otherwise, conventions are the same as in Fig. 7A.

One reason to plot the sensitivity vectors in this way was to see if these different unit populations coded sensory and motor directions using a distributed vector code, as observed in many topographic neural structures (e.g., Sparks, 1989) or if they used a coordinate system, as observed in some brainstem oculomotor structures (e.g., Crawford 1994). Looking at Fig. 7B, one can see that the semiorthogonal and position-opposite$_1$ groups showed no sign of forming a coordinate system but rather showed a fairly even distribution of directions in all of their sensitivity vectors, as predicted by Robinson (1992). The PO_2 group showed somewhat less variability in direction with a slight tendency to align with the standard Cartesian axes, but these were not numerous enough in a given network to make a definite judgment.

In contrast, the vector-propagation group (bottom row) showed a strikingly consistent coordinate system, rotated about the origin ∼45 degrees from the standard Cartesian graph coordinates, in both the retinal-error- and motor-error-sensitivity vectors. This occurred in all trained networks, independent of the random initial weightings. This further served to distinguish this population from the SO group. Thus, across networks, each unit could be assigned to one of the four groups based on consistent characteristics of their sensitivity vectors, with one population—and only one—forming a clear coordinate system.

Although we have focused on the 2-D properties of the unit sensitivity vectors so far, at this point we took the opportunity to look at their torsional components to see how these balanced to give Listing's law. In general, the torsional tuning of motor error in these units (not shown) was somewhat variable, even for the vector-propagation group. Overall, however, torsion varied only slightly from Listing's plane: mean $= -.229$ degrees SD $= .003$ degrees. However, the fact that the torsional motor-error component of these networks was negligible in the actual behavior shows that the positive and negative tilts were perfectly balanced in the overall activation pattern of the units.

These analyses allowed us to characterize individual units as belonging to distinct groups, but they provided only a rough approximation as to the *function* of each group. It was now clear that the network's solution to the reference frame transformation was not completely distributed. Instead, the network developed particular groups of units that processed a particular aspect of the solution. Again, we hypothesized that the vector-propagation group provided the main response to get saccades on target, whereas the main position dependent deviation (that is, convergence toward center) was provided by the position-opposite groups. The role of

the semiorthogonal group was less clear, but it seemed to be the only group that could provide fine tuning such as greater position dependent tilts for greater magnitudes of retinal error.

3.5. Lesion Studies

To test our hypotheses and characterize group function more completely, we conducted *lesion studies*. That is, we removed one or more groups of units from the network's hidden layer and noted the results. First, we conducted control tests in which we randomly lesioned either 25%, 50%, or 75% of units within the networks (not shown). As expected, the results showed a reduction in the magnitude of response with varying degrees of disruption of the position-dependent pattern. These results varied over networks and trials. That is, the magnitude of response was consistently reduced while the amount of disruption to the position-dependent pattern depended on initial position. Thus, no clear pattern emerged from these random lesion controls.

Next we lesioned the vector-propagation group. Recall that we hypothesized that these units represent the main drive to get saccades on target. If this hypothesis was correct, we should see massive degradation of network output in both direction and magnitude of response as a result of such a lesion. Furthermore, if our characterization of the vector-propagation coordinate system is correct, it should be possible to affect one direction selectively without affecting the orthogonal direction.

Figure 10 shows the results of various hidden unit lesions within the vector-propagation group for a particular initial position that was never used during training (25 degrees right and 15 degrees up) and the set of retinal errors that was used with that position. The actual output of the network is shown with thick dark lines, while ideal motor error output is shown with thick light lines. The supplied retinal error is shown in thin dark lines (Fig. 10A only) to highlight the required deviation from a direct retinal-error to motor-error mapping if saccades are to be accurate from all initial positions. Note that after lesioning the network a specific correspondence between network output (dark thick lines) and ideal output (light thick lines) may be difficult since the close correspondence between the ideal pattern and the actual pattern has been disrupted. However, what is important in these lesion studies is not the specific correspondence of each trace in the graph but rather how closely the general pattern of the network response follows that of the ideal response. Also shown in each graph is an inset with directional arrows indicating in which quadrant of the graph the response of vector-propagation units was left *intact.* (Recall that the vector-propagation units use a coordinate system aligned with 45 degree oblique directions).

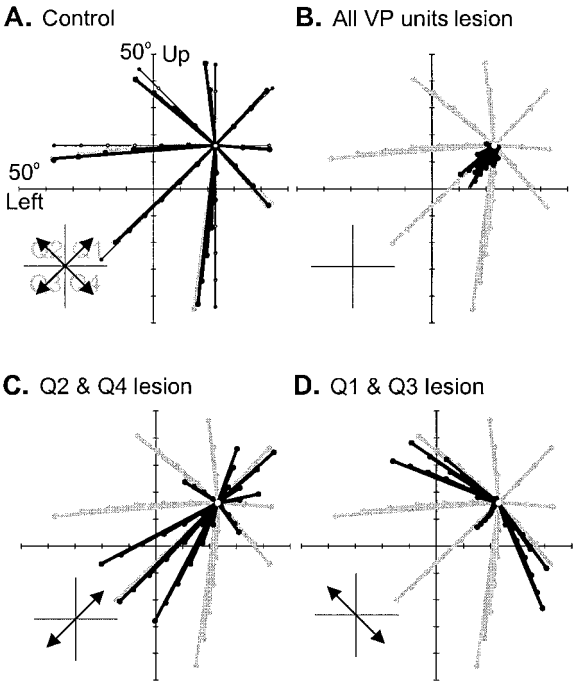

Figure 10. Vector-propagation cell lesions for selected position network 1 (N1). Initial position (25°, 15°). Thick dark lines: Actual network motor error output. Thick light lines: Ideal motor error output. Thin dark lines: Supplied retinal error (plot **A** only). The position has associated with it six retinal errors along each of eight directions where possible without exceeding 50 degree OMR from origin (primary position). Supplied retinal errors are in the same directions as those indicated in Fig. 1B. **A:** Control condition: no lesion. Note that inset axis (all plots) indicates quadrant where vector-propagation units are left *intact*. **B:** All vector-propagation units lesioned. Note complete obliteration of any pattern in response. **C:** Vector-propagation units in Q1 and Q3 lesioned. Note response in the lesioned quadrants are most compromised while those in the other two quadrants are spared. **D:** Vector-propagation units in Q2 and Q4 lesioned.

Figure 10A shows the control condition in which the response of all units within the vector-propagation group (and all other groups) are left intact. Note that, in this plot, the actual output of the network (thick dark lines) closely follows the ideal output (thick light lines) in both direction and magnitude such that a correspondence between the dark lines and the light lines is apparent. Figure 10B shows the effect of lesioning (setting

output to 0) the entire population of vector-propagation units. The result of this lesion was a complete obliteration of the normal pattern, even though ∼50% of hidden-unit output (from semiorthogonal and position-opposite groups) was unaffected. The only remaining effect of these intact groups were weak movements toward center.

Next, we lesioned vector-propagation units whose response lies in the up-left and down-right quadrants. The result is shown in Fig. 10C. Note that both the direction and magnitude components of responses within the quadrant containing the lesion are severely disrupted, whereas output from the network in the up-right/down-left quadrants is less affected, particularly along the intact diagonal coordinate. Similarly, when the opposite coordinates were selectively lesioned (Fig. 10D), saccades in the opposite two quadrants were affected. This confirmed our hypothesis that these units were the main drive to get saccades moving with the right magnitude with approximately the correct direction, as specified componentwise in their intrinsic coordinate system.

To further test this and the role of the remaining groups, we performed the opposite lesion—leaving the vector-propagation group intact while removing other groups. Figure 11 shows the results of these lesions using a different initial position (never used for training: 30 degrees down and 30 degrees right) than that used in our vector-propagation lesions. This position was selected to further highlight the position dependent divergence between retinal error and motor error discussed above. Conventions for actual network output, ideal output, and retinal error are the same as Fig. 10. Figure 11A again shows the control where the output from all groups is intact. Figure 11B shows the network output with only the vector-propagation units intact—that is, position-opposite$_1$, position-opposite$_2$, and semiorthogonal groups did not contribute to the network's performance. Note that direction and magnitude are reasonably intact but there is almost no position-dependent modification of motor output. This again confirmed that the vector-propagation units were the major drive for the correct direction and magnitude of motor error but did not contribute to the required position-dependent modification of motor error.

Figure 11C shows the result of our lesioning the position-opposite$_1$ and position-opposite$_2$ units, leaving both the vector-propagation and semiorthogonal groups intact (we did not attempt to distinguish between the two position-opposite classes because of the very low number of units in the second class and their concomitantly small effect on behavior). Surprisingly, although ∼80% of the units were still functional, we saw a severe compromise of direction and no capacity of the network to perform the position-dependent modification of motor error. Saccade direction diverged wildly from center as if a constant convergence bias had been lost and centrifugal saccade magnitude was greatly enhanced. Note also that this divergence was much stronger for *shorter* saccades such that the pattern was more spread out.

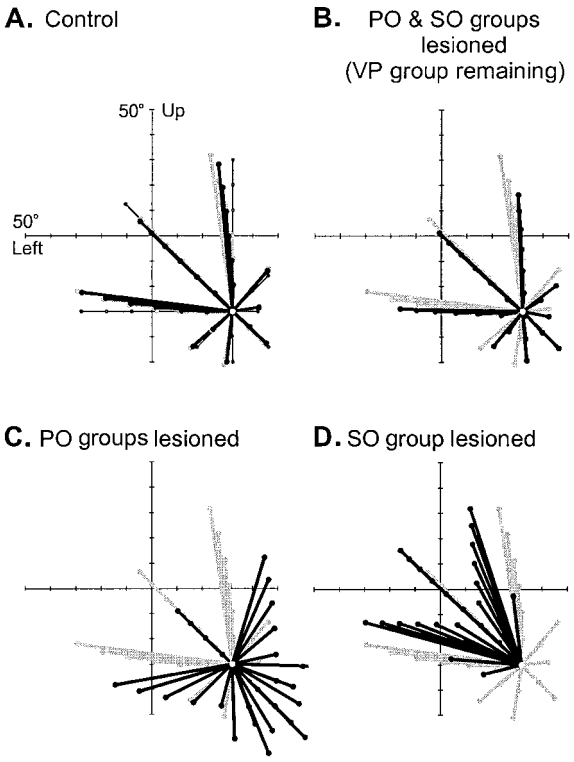

Figure 11. Group lesions. Conventions same as in Fig. 10. **A:** Control: no lesions. **B:** PO$_1$ position-opposite$_1$, PO$_2$ position-opposite$_2$, and SO semiorthogonal lesions (VP vector-propagation group, intact). **C:** PO$_1$ and PO$_2$ lesions. **D:** SO lesions.

When we lesioned only the semiorthogonal group (leaving ∼70% of hidden units intact: Fig. 11D), we saw an equally devastating effect on direction, but this time there was too much convergence. But note that the level of convergence was the same for all saccade magnitudes. Thus surprisingly, a lesion to one of the population groups (semiorthogonal or position-opposite) produced a greater effect than lesioning them both. This led us to conclude that position modulation

was actually being carried out by a *balance* between the contribution of the semiorthogonal and position-opposite groups. Furthermore, whereas the position-opposite group added a constant convergence factor when left intact without the SO group (Fig. 11D), the semiorthogonal group produced less divergence in longer saccades, leading to the phenomenon that longer saccades normally converge more when the two groups work in concert, as required for the ideal geometry (Fig. 11A).

We quantified these observations across networks by testing the networks with the same input set as above, but within each network the PO and SO groups were lesioned (as in Fig. 11B). That is, only the vector-propagation group remained intact. Figure 12 shows the results of this test employing the same conventions used in Fig. 5, except that the *y* axis has a larger range. (◇) shows predictions of a vector-displacement model with no position compensation. Figure 12B is the control condition (primary position) where the vector displacement model (◇) shows no errors, since no position modification is required at that special position.

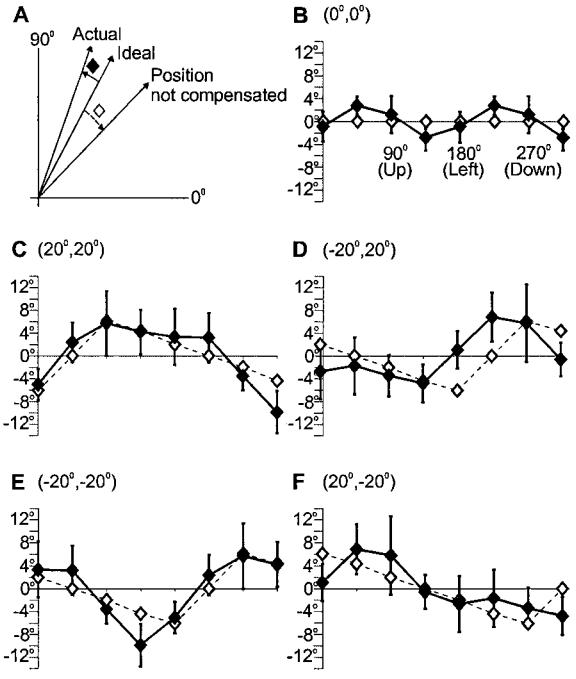

Figure 12. Angular difference measures with vector-propagation lesion across networks. Conventions and initial positions the same as Fig. 5. Note the larger range of the *y* axis. Graph shows that no eye-position compensation takes place without the vector-propagation group (see text for an explanation).

The actual output of the networks (♦) also hovers within ±3 degrees of this response. Looking at Fig. 12C (20°, 20°) we note that errors in the vector displacement model (----◇----) follow the same pattern as in Fig. 5. However, unlike the intact network illustrated in Fig. 5, the actual response of the lesioned networks (—♦—) now more or less followed the incorrect pattern of the vector-displacement model. Thus, in the absence of the semiorthogonal and position-opposite groups, there was no longer any position-dependent modification of response taking place within any of the networks.

4. Discussion

This investigation addressed two main questions. First, can neural networks perform an eye-to-head reference-frame transformation without ever developing head-centric representations of target direction? Second, if so, what type of algorithm is used—a completely distributed transformation or one with recognizable modules?

4.1. The Oculomotor Reference-Frame Transformation

Most theoretical studies of oculomotor reference frame transformations have focused on the problem of remembering target locations across saccades. Traditionally, these have assumed that the system develops a head-centric map of visual space to remember targets independently of subsequent saccades (Soechting and Flanders, 1992; Zee et al., 1976; Howard, 1982). Even the original *implicit gain-field* model of Zipser and Andersen (1988) employed such a map in the model's output layer (although the authors suggested that this might be avoided in the real system). Unfortunately, this is at odds with the known physiology of the saccade generator, which seems to rely largely on displacement signals (Moschovakis and Highstein, 1994; Colby and Goldberg, 1999). Moreover, recent studies have shown that target directions can be remembered across saccades without a headcentric map by "remapping" their representations retinotopically (Bozis and Moschovakis, 1998; Duhamel et al., 1992; Goldberg and Bruce, 1990; Krommenhoek et al., 1993; Waitzman et al., 1991; Henriques et al., 1998; Batista et al., 1999). Be this as it may, the eye-to-head reference-frame transformation problem does not go away so easily.

The oculomotor system would still require an eye-to-head reference-frame transformation—downstream from the remapping mechanisms—to deal with the 3-D geometry of saccade execution. For saccades to be accurate and obey Listing's law, 2-D oculocentric retinal saccade signals must be correctly transformed into 3-D head-centric motor-error commands in Listing's plane (Crawford and Guitton, 1997; Klier and Crawford, 1998). Crawford and Guitton (1997) modeled this with the use of explicit representations of gaze direction and 3-D eye position in head coordinates to perform this transformation. But this still conflicts with the known physiology. Therefore, we asked if a neural network model trained to perform the same task could provide an alternative solution.

The neural network in the current study clearly learned the correct position dependent modification of retinal error to provide the reference-frame transformation for accurate motor-error commands, but it did so without an explicit representation of desired 3-D eye orientation or target direction in head coordinates. That is, no one unit coded for a specific spatial location or direction of retinal error. This may not answer the question of whether the brain actually uses head-centric representations of target direction or desired eye orientation, but it does answer the question of whether it must do so to perform an eye-to-head reference-frame transformation. The answer to that question is a definite no. This is not to say that other nonsaccadic systems do not use head-centric maps (Duhamel et al., 1997), but it goes a long way to explaining why they are not necessary in the saccade generator (Colby and Goldberg, 1999).

4.2. Modularity and Coordinate Systems

There has been a general consensus, at least in the visuomotor field, that artificial neural networks do not employ recognizable algorithms or representations (Robinson, 1992; Stein et al., 1992) but rather produce highly distributed solutions. However, this may have been a product of the network analysis used.

In our network, a physiological type of approach proved to be more fruitful than a statistical analysis of relationships between input, output, and program weights. Remarkably, this showed that individual hidden units within the network did have understandable functions. The solution to the reference-frame transformation problem was not completely distributed throughout the network, but neither was it solved using a unitary multiplication algorithm as proposed by Crawford et al. (2000) for similar transformations in the cortex. Rather, it was solved by parallel groups, or modules, where each of these parallel modules performed a specific subtask.

Our analysis suggested that the dominant module (vector-propagation) simply performed the first-order linear mapping between retinal and motor error, whereas higher order nonlinearities were handled by the other modules. The second-order convergence effect (Figs. 3, 4, and 6) was handled by the actions of the position-opposite groups, which mainly directed motor error in the direction opposite to initial position. However, this was insufficient because the amount of convergence also depended on retinal error magnitude and direction (e.g., Fig. 4). Thus, a third major group (semiorthogonal) was required to modulate the effects of the position-opposite groups through a system of balance. Clearly, more than one unit was required to implement each of these groups, which accounted for the disproportionately low number of functional groups compared with the number of hidden units.

Among these groups, only the vector-propagation group consistently formed the same orthogonal coordinate system in each network, despite random initial settings. Moreover, this coordinate system aligned with Listing's plane, as observed in the coordinates of the real saccade generator (Crawford, 1994; Crawford and Vilis, 1992). Although the mathematical reason for the development of this coordinate system is unclear at this time, this result is difficult to trivialize because this diagonal coordinate system differed from the Cartesian coordinates used to input data and was not observed in the other unit types. Moreover, since coordinates are not required to specify direction (Soechting and Flanders, 1992; Sparks, 1989; Georgopulos et al., 1982), they must have conferred some specific advantage—such as economy of directional representation among a small number of units.

In any case, although the overall reference frame transformation was implicit across these parallel modules, the network did employ an identifiable algorithm and explicit coordinate system to do so, contrary to general expectation (Robinson, 1992). Clearly, developmental rules also contribute to the organization of the saccade generator, but this study shows that error-driven training alone is sufficient to give rise to a considerable amount of functional modularity.

4.3. Implications for Oculomotor Physiology

Where then might a neurophysiologist look for the functional modules observed in our network's hidden-unit layer? To begin with, we need to consider how the input and output layers of our network might correspond to functional neuroanatomy. First the outputs. If oculomotor short-lead burst neurons encode the derivative of eye orientation rather than angular velocity—as suggested in several recent studies (Crawford and Guitton 1997; Quaia and Optican, 1998; Hepp et al., 1999)—then the motor-error output of our network would be the appropriate signal to drive these neurons. Thus, we would interpret the output layer of our network as representing the total ensemble of inputs to short-lead burst neurons, which would include inputs from the superior colliculus and other sources (e.g., Quaia et al., 1999; Helmchen et al., 1996).

The visual signal starts out coded in oculocentric coordinates—as in our input layer, albeit represented in a more distributed topographic form. But how far downstream does such a code persist? One clue is that the pattern of converging saccades that we simulated for stimulation of sites coding a fixed oculocentric retinal error (Fig. 6) resembles the pattern of saccades evoked by stimulation of various cortical and subcortical sites (Schlag and Schlag-Rey, 1987; Bon and Luchetti, 1992; Schall, 1991; Freedman et al., 1996; Russo and Bruce, 1993). In particular, in the study by Russo and Bruce (1993), it was shown that saccades induced by electrical stimulation of frontal eye fields and supplementary eye fields in macaque monkeys showed an orbital perturbation index (defined in the Results section) ranging between $\sim$0.0 to $-$0.5 with a mean K_h of -0.13 (SD $= 0.11$) for SEF site stimulations and -0.16 (SD $= 0.11$) for FEF site stimulations. These values are comparable to those that were observed with our model, and close inspection of their data also reveals that the convergence index increased with saccade size—as in our model. When observed experimentally, such convergence patterns have generally been interpreted to signify some artifact (that is, error) in the downstream structures that account for eye position (e.g., Russo and Bruce, 1993). However, the current analysis suggests that orbital convergence like that reported by Russo and Bruce could mean that these sites are simply coding true retinal error, with the *correct* geometric transformations occurring downstream.

If so, could this retinal code persist as far as the superior colliculus? In their 1991 study of the colliculus, Van Opstal et al. did not distinguish between the kind of retinal-error and motor-error signals that were used in our model, so this question remains open. Moreover, in another study (Van Opstal et al., 1995), they showed that the superior colliculus possesses eye-position-dependent "gain fields" that in theory could be used to construct many different types of spatial codes. However, the head-free colliculus stimulation data of Freedman et al. (1996) shows a pattern of convergence that could be consistent with a retinal code. Furthermore, preliminary results of a very recent study designed to test this specificially by computing the 3-D geometry of such movements support this conclusion (Klier et al., 2000). Therefore, our current thinking is that the input layer of our neural network could correspond to the output code of the superior colliculus—broken down into its horizontal and vertical components—with our hidden layer intervening (synaptically) between this and the short-lead burst neurons.

Before testing this specific scheme, it is important to note that we simplified our feed-forward model considerably to facilitate analysis and that the inputs to the model were simplified vector components of retinal error and eye position. The real brainstem saccade generator, besides having many more neurons, also faces other problems related to the spatiotemporal mapping from the retinotopic map of the colliculus to reticular formation burst neurons (Van Opstal and Van Gisbergen, 1981; Tweed and Vilis, 1990b; Quaia and Optican, 1997). Among these are temporal dynamics (Van Gisbergen et al., 1981; Munoz and Wurtz, 1995), dealing with deviations of torsion from Listing's plane (Crawford and Vilis, 1991; Tweed et al., 1998; Crawford et al., 1999), and using additional lateral connections and feedback control (Quaia et al., 1998). However, our network shares one important characteristic with the real brain (that is, the connections between the superior colliculus and burst neurons) in that it had to solve the reference frame problem within approximately two synapses, perhaps aided by certain parallel streams.

For this reason, the networks' method of dividing the job into parallel task modules may reflect a realistic physiological strategy. It therefore seems reasonable that the brainstem would employ a main channel resembling the vector-propagation units in our network (units with similar sensory and motor on-directions that develop a coordinate system—like the progression seen from colliculur burst neurons to short lead

burst neurons) to do the "main job" of getting the saccade going in the general direction of retinal error. Likewise, the position-dependent modulations that we observed in our position-opposite and semiorthogonal units could likely be implemented in some side-path within the reticular formation or perhaps the cerebellum (Vilis and Hore, 1981; Russo and Bruce, 1993; Quaia et al., 1998). Investigators interested in such a process should look for cells with eye-position sensitivity opposite to its saccade tuning (position-opposite type) and cells with visual, eye position, and saccade sensitivity in pseudorandom, nonparallel arrangements (semiorthogonal). Moreover, if the real brain employs the same system of balance between such cells, then it may be possible for experimental or clinical lesions to uncouple this balance (Fig. 11), with the potentially devastating effects that were simulated in the results section. But now that the current study has established a basic theoretical framework, it will also be important to repeat these simulations in a more constrained network designed to more closely emulate the properties of neurons in the superior colliculus and brainstem.

4.4. General Implications

This study has two important general implications. The first is that visuomotor transformations can accomplish the reference-frame transformations required for accurate motor behavior without the need for intermediate spatial maps. Although we demonstrated this for the oculomotor system, the same principle applies to, for example, the arm-control system, where visual signals must be transformed from a retinal frame in the parietal cortex (Batista et al., 1999) and premotor cortex (Mushiake et al., 1997) to body-centric vector codes in the motor cortex (Georgopulos et al., 1982). Our network demonstrates that this can be done simply by modulating vector codes by position inputs. However, contrary to one suggestion (Crawford et al., 2000), our network suggests that some of these modulations may be performed by parallel modules rather than entirely intrinsic network computations.

Second, this study shows that neural networks can develop task-related modules like those observed in the real brain, simply through the use of a training algorithm, without the need of additional developmental or anatomic constraints. It should be noted that the *functional* modules discovered in our network differ from the *spatial* clustering of units with similar temporal spiking patterns found in some artificial neural network models (Parodi et al., 1998; Draye et al., 1997; Xing et al., 1996). Closer results to the functional organization of hidden units found in this investigation may be models involving the organization of visual cortex cells (Ernst et al., 1999; Bell and Sejnowski, 1997).

Currently, it is unclear whether the apparent lack of reports of such specialization in many previous network models is due to a fundamental difference (such as our simplified vector inputs of retinal error and eye position), or whether this discrepancy is due to differences in the way that the network models were analyzed. Thus, our next step is to investigate whether a similar model will produce such functional specialization while using more realistic input values. Nevertheless, the current approach provides a more optimistic view than popular expectation for analysing real neural networks and for further theoretical studies of genetically programmed anatomy and learning, in determining function modularity.

Appendix

The following gives a summary of the mathematical formulas used in designing the networks in this study.

Transfer Function

Each of the three layers (input, hidden, and output) of the network used a sigmoid transfer function given by

$$f(x) = \frac{\gamma}{1 + \exp(-\sigma x)} - \eta, \tag{1}$$

where γ is range of the transfer function σ is slope parameter (steepness of the transfer function), η is maximum value of the transfer function.

The derivative of the transfer function was used in back-propagating error through the network:

$$f'(x) = \frac{\sigma}{\gamma}[\eta + f(x)][\gamma - \eta - f(x)]. \tag{2}$$

Error Term

The error term for each unit in the output layer was calculated using

$$err_i = (t_i - z\_out_k)\, f'(z\_in_i), \tag{3}$$

where i is output unit, t is target vector, *z_out* is output layer output, *z_in* is output layer input.

Each unit in the hidden layer calculated the change term using the error term from the output layer multiplied by the current hidden layer weight for that unit:

$$c\_in_j = \sum_{k=1}^{m} err_k w_{jk}, \tag{4}$$

where j is unit j in hidden layer, k is unit k in output layer, *c_in* is change term, err_k is error term from output unit k, w_{jk} is weight between hidden unit j and output unit k.

The error term for the hidden layer is given by

$$err_j = c\_in_j \, f'(y\_in_j), \tag{5}$$

where j is unit j in hidden layer, *c_in* is change term, *y_in* is input of hidden unit.

Weight Update with Momentum

The weight correction term (output layer shown) was calculated by

$$\Delta w_{jk} = \alpha \delta_k y_j, \tag{6}$$

where α is learning rate, δ is error term, y is output of hidden unit.

A momentum term (10%) was added to the weight update procedure:

$$\Delta w_{jk}(t+1) = \alpha \Delta_k y_j + \mu \delta w_{jk}(t), \tag{7}$$

where the change in weight w_{jk} at time $(t+1)$ is the sum of the weight correction term (#5) and the product of the momentum term (μ) with change in weight w_{jk} at time t.

Acknowledgments

We thank D.B. Tweed for comments. This work was supported by a Canadian Natural Sciences and Engineering Research Council grant to J.D. Crawford. J.D. Crawford is supported by a Canadian Medical Research Council Scholarship. M.A. Smith is supported by a Canadian Natural Sciences and Engineering Research Council scholarship.

References

Batista AP, Buneo CA, Snyder LH, Andersen RA (1999) Reach plans in eye-centered coordinates. *Science* 285:257–260.

Bell AJ, Sejnowski TJ (1997) The "independent components" of natural scenes are edge filters. *Vision Res.* 37:3327–3338.

Bon L, Lucchetti C (1992) The dorsomedial frontal cortex of the macaca monkey: Fixation and saccade-related activity. *Exp. Brain Res.* 89:571–580.

Bozis A, Moschovakis AK (1998) Neural network simulations of the primate oculomotor system. III. A one-dimensional, one-directional model of the superior colliculus. *Biol. Cybern.* 79: 215–230.

Colby CL, Duhamel JR, Goldberg ME (1995) Oculocentric spatial representation in parietal cortex. *Cereb. Cortex.* 5:470–481.

Colby CL, Goldberg ME (1999) Space and attention in parietal cortex. *Ann. Rev. Neurosci.* 22:319–349.

Crawford JD (1994) The oculomotor neural integrator uses a behavior-related coordinate system. *J. Neurosci.* 69:11–23.

Crawford JD, Ceylan MZ, Klier EM, Guitton D (1999) Three-dimensional eye-head coordination during gaze saccades in the primate. *J. Neurophysiol.* 81:1760–1782.

Crawford JD, Guitton D (1997) Visual motor transformations required for accurate and kinematically correct saccades. *J. Neurophysiol.* 78:1447–1467.

Crawford JD, Henriques DYP, Vilis T (2000). Curvature of visual space under vertical eye rotation: Implications for spatial vision and visuomotor control. *J. Neurosci.* 20:2360–2368.

Crawford JD, Vilis T (1991) Axes of eye rotation and Listing's law during rotations of the head. *J. Neurophysiol.* 65: 407–423.

Crawford JD, Vilis T (1992) Symmetry of oculomotor burst neuron coordinates about Listing's plane. *J. Neurophysiol.* 68:432–448.

Crawford JD, Vilis T (1993) Modularity and parallel processing in the oculomotor integrator. *Exp. Brain. Res.* 96:443–456.

Demer JL, Miller JM, Poukens V, Vinters HV, Glasgow BJ (1995) Evidence for fibromuscular pulleys of the recti extraocular muscles. *Invest. Ophthalmol. Vis. Sci.* 36:1125–1136.

Draye JP, Cheron G, Libert G, Godaux E (1997) Emergence of clusters in the hidden layer of a dynamic recurrent neural network. *Biol. Cybern.* 76:365–374.

Duhamel J, Bremmer F, BenHamed S, Graf W (1997) Spatial invariance of visual receptive fields in parietal cortex neurons. *Nature* 389:845–848.

Duhamel JR, Colby CL, Goldberg ME (1992) The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255:90–92.

Ernst U, Pawelzik K, Tsodyks M, Sejnowski TJ (1999) Relation between retinotopic and orientation maps in visual cortex. *Neural Comput.* 11:375–379.

Flanders M, Daghestani L, Berthoz A (1999) Reaching beyond reach. *Exp. Brain Res.* 126:19–30.

Freedman EG, Sparks DL (1997) Activity of cells in the deeper layers of the superior colliculus of Rhesus monkey: Evidence for a gaze displacement command. *J. Neurophysiol.* 78:1669–1690.

Freedman EG, Standford TR, Sparks DL (1996) Combined eye-head gaze shifts produced by electrical stimulation of the superior colliculus in rhesus monkeys. *J. Neurophysiol.* 76:927–951.

Freedman JA, Skapura DM (1991) *Neural Networks: Algorithms, Applications, and Programming Techniques.* Addison Wesley, Toronto.

Fuchs AF, Kaneko CRS, Scudder CA (1985) Brainstem control of saccadic eye movements. *Ann. Rev. Neurosci.* 8:307–337.

Georgopulos AP, Kalaska JF, Camaniti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* 2:1527–1537.

Goldberg ME, Bruce CJ (1990) Primate frontal eye fields. III. Maintenance of a spatially accurate saccade signal. *J. Neurophysiol.* 64:489–508.

Grossberg S, Kuperstein M (1986) *Neural Dynamics of Adaptive Sensory Motor Control, Ballistic Eye Movements.* Elsevier North-Holland, North-Holland.

Helmchen C, Rambold H, Buttner U (1996) Saccade-related burst neurons with torsional and vertical on-directions in the interstitial nucleus of Cajal in the alert monkey. *Exp. Brain. Res.* 112:63–78.

Henn V, Hepp K, Vilis T (1989) Rapid eye movement generation in the primate: Physiology, pathophysiology and clinical implication. *Revue Neurologique* (Paris) 145:540–545.

Henriques DYP, Klier EM, Smith MA, Lowy D, Crawford JD (1998) Gaze-centered remapping of remembered visual space in an open loop-pointing task. *J. Neurosci.* 18:1583–1594.

Hepp K, Cabungcal JH, Duersteler M, Hess BJM, Scherberer H, Straumann D, Suzuki Y, Van Opstal AJ, Henn, V (1999) 3D structure of the reticular saccade generator in the monkey (Abstract). 29th Annual Meeting of the Society for Neuroscience, Abstract #661.6 (book 2).

Hepp K, Van Opstal AJ, Straumann D, Hess BJ, Henn V (1993) Monkey superior colliculus represents rapid eye movements in a two-dimensional motor map. *J. Neurophysiol.* 69:965–979.

Howard IP (1982) *Human Visual Orientation.*Wiley, New York.

Jürgens R, Becker W, Kornhuber HH (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. *Bio. Cybern.* 39:87–96.

Klier EM, Crawford JD (1998) Human oculomotor system accounts for 3-D eye orientation in the visual-motor transformation for saccades. *J. Neurophysiol.* 80:2274–2294.

Klier EM, Wang H, Crawford JD (2000) Stimulation of the interstitial nucleus of Cajal produces torsional and vertical head rotations (Abstract). *Society for Neurosci. Abstracts* 25:660.18.

Krommenhoek KP, Van Opstal AJ, Gielen CC, Van Gisbergen JA (1993) Remapping of neural activity in the motor colliculus: A neural network study. *Vision Res.* 33:1287–1298.

Krommenhoek KP, Wiegerinck WA (1998) A neural network study of precollicular saccadic averaging. *Biol. Cybern.* 78:465–477.

Liu L, Sun H, Guo A (1997) Transformation of sensory signals into commands for saccadic eye movements: A neural network study. *J. Theor. Biol.* 189:121–131.

Moschovakis AK, Highstein SM (1994) The anatomy and physiology of primate neurons that control rapid eye movements. *Ann. Rev. Neurosci.* 17:465–488.

Munoz DP, Wurtz RH (1995) Saccade-related activity in monkey superior colliculus II. Spread of activity during saccades. *J. Neurophysiol.* 73:2234–2348.

Mushiake H, Tanatsugu Y, Tanji J (1997) Neuronal activity in the ventral part of premotor cortex during target-reach movement is modulated by direction of gaze. *J. Neurophysiol.* 78:567–567.

Optican LM, Miles FA (1985) Visually induced adaptive changes in primate saccadic oculomotor signals. *J. Neurophysiol.* 54: 940–958.

Parodi P, Jimbo Y, Kawana A, Macri D, Torre V (1998) Segmentation of the response of a neuronal network into clusters with similar activity. *BioSystems.* 48:171–178.

Quaia C, Aizawa H, Optican LM, Wurtz RH (1998) Reversible inactivation of monkey superior colliculus. II. Maps of saccadic deficits. *J. Neurophysiol.* 79:2097–2110.

Quaia C, Lefevre P, Optican LM (1999) Model of the control of saccades by superior colliculus and cerebellum. *J. Neurophysiol.* 82:999–1018.

Quaia C, Optican LM (1997) Model with distributed vectorial premotor bursters accounts for the component stretching of oblique saccades. *J. Neurophysiol.* 78:1120–1134.

Quaia C, Optican LM (1998) Commutative saccadic generator is sufficient to control a 3-D ocular plant with pulleys. *J. Neurophysiol.* 6:3197–3215.

Raphan T (1998) Modelling control of eye orientation in three dimensions. I. Role of muscle pulleys in determining saccadic trajectory. *J. Neurophysiol.* 79:2653–2667.

Robinson DA (1981) Control of eye movements. In: Brooks VB, ed. *Handbook of Physiology: The Nervous System.* Visual motor control. Bethesda, MD. *Am. Physiol. Soc.* 2:1275–1320.

Robinson DA (1992) Implications of neural networks for how we think about brain function. *Behav. and Brain Sci.* 15:644–655.

Rumelhart DE, Hinton GE, Williams RJ (1986) Learning internal representations by error propagation. In: Rumelhart DE, McClelland JL, eds. *Parallel Distributed Processing.* MIT Press, Cambridge, Mass. vol. 1, pp. 318–362.

Russo GS, Bruce CJ (1993) Effect of eye position within the orbit on electrically elicited saccadic eye movements: A comparison of the Macaque monkey's frontal and supplementary eye fields. *J. Neurophysiol.* 69:800–818.

Schall JD (1991) Neuronal activity related to visually guided saccadic eye movements in the supplementary motor area of rhesus monkeys. *J. Neurophysiol.* 66:530–558.

Scherberger H, Cabungeal JH, Heep K, Suzuki Y, Straumann D, Henn V (1998) Pre-saccadic reticular burst neuron on-directions rotate with ocular counterroll in monkeys. *Society for Neuroscience Abstracts* 24:60–65.

Schlag J, Schlag-Rey M (1987) Evidence for a supplementary eye field. *J. Neurophysiol.* 57:179–200.

Snyder LH, Batista AP, Andersen RA (1998) Change in motor plan, without a change in the spatial locus of attention, modulates activity in the posterior parietal cortex. *J. Neurophysiol.* 79:2814–2819.

Soechting JF, Flanders M (1989) Sensorimotor representations for pointing to targets in three-dimensional space. *J. Neurophysiol.* 62:582–594.

Soechting JF, Flanders M (1992) Moving in three-dimensional space: Frames of reference, vectors, and coordinate systems. *Ann. Rev. Neurosci.* 15:167–191.

Sparks DL (1989) The neural encoding of the location of targets for saccadic eye movements. *J. Exp. Biol.* 146:195–207.

Stein JF (1992) The representation of egocentric space in the posterior parietal cortex. *Behav. and Brain. Sci.* 15:691–700.

Tweed D, Haslwanter T, Fetter M (1998) Optimizing gaze control in three dimensions. *Science* 281:1363–1366.

Tweed DB, Vilis T (1990a) Geometric relations of eye position and velocity vectors during saccades. *Vision Res.* 30:111–127.

Tweed DB, Vilis T (1990b) The superior colliculus and spatiotemporal translation in the saccadic system. *Neural Networks* 3:75–86.

Van Gisbergen JA, Robinson DA, Gielen S (1981) A quantitative

analysis of generation of saccadic eye movements by burst neurons. *J. Neurophysiol.* 3:417–442.

Van Opstal AJ, Hepp, K (1995) A novel interpretation for the collicular role in saccade generation. *Biol. Cybern.* 73:431–445.

Van Opstal AJ, Hepp K, Hess BJ, Straumann D, Henn V (1991) Two rather than three-dimensional representation of saccades in monkey superior colliculus. *Science* 252:1313–1315.

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

Vilis T, Hore J (1981) Characteristics of saccadic dysmetria in monkeys during reversible lesions of medial cerebellar nuclei. *J. Neurophysiol.* 46:828–838.

Von Helmholtz H (1925) *Treaties on Physiological Optics* (vol. 3). Translated by J.P.C. Southall. Optical Society of America, Rochester, NY.

Waitzman DM, Ma TP, Optican LM, Wurtz RH (1991) Superior colliculus neurons mediate the dynamic characteristics of saccades. *J. Neurophysiol*. 66:1716–1737.

Woodwoth RS (1899) The accuracy of voluntary movement. *Psychol. Rev. Monogr. Suppl. 3*.

Xing J, Gerstein GL (1996) Networks with lateral connectivity. II. Development of neuronal grouping and corresponding receptive field changes. *J. Neurophysiol.* 75:200–216.

Zee DS, Optican LM, Cook JD, Robinson DA, Engel WK (1976) Slow saccades in spinocerebellar degeneration. *Arch. Neurol.* 33:243–251.

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