Influence of Saccade Efference Copy on the Spatiotemporal Properties of Remapping: A Neural Network Study

Gerald P. Keith,[1,2] Gunnar Blohm,[4] and J. Douglas Crawford[1,2,3]

[1]*Centre for Vision Research and Canadian Institute of Health Research Group,* [2]*Departments of Psychology and* [3]*Biology and Kinesiology and Health Sciences York University, Toronto; and* [4]*Department of Computational Neuroscience Queen's University, Kingston, Ontario, Canada*

Keith GP, Blohm G, Crawford JD. Influence of saccade efference copy on the spatiotemporal properties of remapping: a neural network study. *J Neurophysiol* 103: 117–139, 2010. First published October 21, 2009; doi:10.1152/jn.91191.2008. Remapping of gaze-centered target-position signals across saccades has been observed in the superior colliculus and several cortical areas. It is generally assumed that this remapping is driven by saccade-related signals. What is not known is how the different potential forms of this signal (i.e., visual, visuomotor, or motor) might influence this remapping. We trained a three-layer recurrent neural network to update target position (represented as a "hill" of activity in a gaze-centered topographic map) across saccades, using discrete time steps and backpropagation-through-time algorithm. Updating was driven by an efference copy of one of three saccade-related signals: a transient visual response to the saccade-target in two-dimensional (2-D) topographic coordinates (*Vtop*), a temporally extended motor burst in 2-D topographic coordinates (*Mtop*), or a 3-D eye velocity signal in brain stem coordinates (*EV*). The Vtop model produced presaccadic remapping in the output layer, with a "jumping hill" of activity and intrasaccadic suppression. The Mtop model also produced presaccadic remapping with a dispersed moving hill of activity that closely reproduced the quantitative results of Sommer and Wurtz. The EV model produced a coherent moving hill of activity but failed to produce presaccadic remapping. When eye velocity and a topographic (Vtop or Mtop) updater signal were used together, the remapping relied primarily on the topographic signal. An analysis of the hidden layer activity revealed that the transient remapping was highly dispersed across hidden-layer units in both Vtop and Mtop models but tightly clustered in the EV model. These results show that the nature of the updater signal influences both the mechanism and final dynamics of remapping. Taken together with the currently known physiology, our simulations suggest that different brain areas might rely on different signals and mechanisms for updating that should be further distinguishable through currently available single- and multiunit recording paradigms.

INTRODUCTION

Behavioral studies have shown that humans and primates are able to maintain an accurate sense of remembered target positions across saccades (Baker et al. 2003; Hallett and Lightstone 1976; Matin et al. 1969; Mays and Sparks 1980; Miller 1980; Schiller and Sandel 1983; Schlag-Rey et al. 1989; Sparks and Porter 1983). Behavioral (Cai et al. 1997; Henriques et al. 1998; Honda 1989; Matin et al. 1969; Ross et al. 1997) and neurophysiological (Colby et al. 1995; Russo and Bruce 1996; Sparks 1988, 1989) evidence suggests that target positions are stored in gaze-centered coordinates. Such posi-

tions must be spatially updated across gaze shifts. Brain areas associated with saccade programming, including the superior colliculus (SC), frontal eye-fields (FEF), and lateral intraparietal cortex (LIP), show activity remapping across saccades that has been associated with this updating (Duhamel et al. 1992; Mays and Sparks 1980; Umeno and Goldberg 1997; Walker et al. 1995). What remains unsolved is: what signals drive this remapping and how might these signals influence activity *during* the remapping?

The observation of remapped activity that for some neurons precedes saccade onset (Duhamel et al. 1992; Umeno and Goldberg 1997; Walker et al. 1995) suggests that saccade efference copy signals drive remapping. There are several such candidate signals. The superficial and deeper layers of the SC show a transient visual burst in response to the appearance of a saccade target (Bruce and Goldberg 1985; Mays and Sparks 1980), and the deeper layers of the SC also show temporally extended, saccade-related motor bursts (Munoz and Wurtz 1995). These signals use a two-dimensional (2-D) topographic place code (Klier et al. 2003; van Opstal et al. 1991) that could drive remapping. As well, the necessary saccade-metric information exists at the level of brain stem burst neurons that encode saccade velocity using a 3-D rate coding scheme (Crawford and Vilis 1992; Luschei and Fuchs 1972; Suzuki et al. 1995). Recent experiments involving inactivation of thalamic nuclei show that the SC provides at least part of the signal that drives remapping (Sommer and Wurtz 2006), but remapping has never been entirely obliterated, so it is possible that other signals are used as well.

There is no clear consensus in the theoretical literature on what type of efference copy signals might be used to drive updating. For example, the vector subtraction model of Quaia et al. (1998) used the saccade-related burst signal in the frontal eye field, that of Medendorp et al. (2003) used efferent copies of brain stem eye position and velocity information, and Droulez and Berthoz (1991) used eye velocity alone. Earlier visual signals might have the advantage of prediction, whereas late velocity signals might have the advantage of providing a more accurate measure of the actual movement. Moreover, it is possible, and perhaps likely, that different combinations of signals are used in different brain areas. But currently there is no rigorous theoretical framework for physiologists to distinguish between these possibilities.

One way to distinguish between these models is to explore how the spatiotemporal form of the efference copy updater signal affects the progression of the remapping associated with updating. Two contrasting hypotheses for how remapping may

Address for reprint requests and other correspondence: J. D. Crawford, York Centre for Vision Research, Computer Science Building, York University, 4700 Keele St., Toronto, Ontario M3J 1P3, Canada (E-mail: jdc@yorku.ca).

occur envisage a "*moving* (or *spreading*) *hill*" or a "*jumping hill*" of activity.[1] In the most comprehensive study of the remapping process to date, Sommer and Wurtz (2006) recorded the evolution in activity during saccades for FEF neurons, the receptive fields of which were centered at the *remapping midpoint* (i.e., the midpoint between initial and updated target position). They found that there was no increase in activity at this midpoint during remapping. This was consistent with the view that the hill of activity associated with this target position jumped rather than moved or spread. However, the pattern of remapping need not be constrained to these two options. For example, Walker et al. (1995) found that the presaccadic remapped activity in the SC was suppressed during saccades, reappearing only after the saccade. Also, there is emerging evidence (e.g., Bremmer et al. 2009; Kubischik and Bremmer 1999) that remapping follows different patterns in different brain areas.

In analyzing these extremely difficult experiments, it would be helpful to know what possible forms beyond the moving hill/jumping hill dichotomy can occur during remapping, and updater signals are associated with each. Developing a formal theoretical framework was the purpose of the current study.

Rather than using *designed* or *constructed* updating neural network models, where updating is achieved by imposing very specific recurrent connectivities with tight spatial tuning on a set of topographically arranged neurons (Bozis and Moschovakis 1998; Droulez and Berthoz 1991; Quaia et al. 1998), we chose the approach of *trained* updating neural network models (Mitchell and Zipser 2001; White and Snyder 2004; Xing and Andersen 2000), networks with universal feed-forward connectivity between adjacent network layers, and recurrent connectivity within the hidden network layer. The weights of these connections are learned by means of a backpropagation algorithm as the network is trained to perform the updating task. The use of trained models has the advantage of producing emergent (and often unexpected) properties that arise from the training rather than being explicitly input by the investigator. If these emergent properties resemble observations made in experimental neuroscience, the underlying structure of the artificial network may be viewed as providing clues to the underlying mechanisms of the real neural network (e.g., Zipser and Andersen 1988) as well as demonstrating general principles.

Previous neural network studies of updating used 1-D (White and Snyder 2004) or 2-D (Mitchel and Zipser 2001; Xing and Andersen 2000) geometry and did not focus on the evolution of activity in the network during remapping. In our simulations, we have used the full 3-D geometry of updating to place constraints on updating that influence the type of signals required (Medendorp et al. 2002; Smith and Crawford 2005). In our previous studies (Keith and Crawford 2008; Keith et al. 2007), we trained a feed-forward neural network to perform remapping in a single time step and then analyzed the hidden-layer units, discovering in the process that different mechanisms were used for the "linear" and "nonlinear" aspects of remapping. In our current study, we used a similar network architecture but added an explicit representation of time (in distinct time steps) and recurrent connections by means of which the information contained in the network was sustained across time. Updating networks were trained using one of three efference copy signals to drive the updating: the early, transient visual activity of the superficial SC (*Vtop*); the later, sustained motor activity of the deeper SC (*Mtop*)—both of these being 2-D topographic signals; and finally the feedback eye velocity burst activity (*EV*) in 3-D brain stem coordinates. We then analyzed the emergent properties of both the hidden and output layer to determine the influence of the input updater signal on both the mechanism and intrasaccadic dynamics of the resulting updating signals. Our results show different patterns of activity during remapping in both the hidden and output layers of the network and yield several predictions that can be tested in the brain by physiologists.

METHODS

The goal of the current study was to examine how the spatiotemporal properties of signals used to drive the updating of target position across saccades (*updater signals*) affect the remapping of neural activity associated with this updating. We trained three-layer,[2] recurrent neural networks using different updating models (i.e., different combinations of updater signals) to perform spatial updating across an intervening saccade. The updating considered was that associated with the double-saccade task (Hallett and Lightstone 1976). Because it is generally believed that target positions are stored in retinal coordinates, the remembered position of the second target in this task must be updated across the saccade to the first target. The updater signals used were efference copies of signals associated with the first saccade.

To examine the manner in which the remapping occurs, we explicitly represented the temporal evolution of both target position and updater signals in our networks, using discrete time steps. This approach differed from that of our earlier study, which focused on the geometric rather than temporal aspects of remapping, and did not vary the form of the updater signal (Keith and Crawford 2008).

Standard network architecture

The architecture of our *standard updating neural network* is shown in Fig. 1*A*. The basic architecture of this network, in terms of inputs and outputs, was similar to that of our previous study (Keith and Crawford 2008). A brief description of the standard network is given here. There were also two variations of this network, one using a warped topographic field, another using shorter time steps. Details of all networks are given in the supplementary materials.[3]

The first, input layer of the network represented the signals observed in various brain structures that were used in the updating. We postulated that copies of these signals were sent to a common brain area, here represented by the hidden layer of the network, where the updating computation was performed. The initial second-target position input signal (TP_2) signal was encoded in the input layer as a hill of activity in a 2-D retinotopic array of units with Gaussian receptive fields. Also encoded in the input layer were the updater signals. Two principal types of updater signals were considered in our updating models—place-encoded signals: visual and motor topographic (Vtop and Mtop) target-position signals seen in the FEF and SC; and rate-encoded signals: eye velocity (EV) and position (EP) signals of the reticular nuclei. These signals are described in the next section.

---

[1] These two models were first proposed not for remapping itself but for the dynamic control of saccades in the SC (Keller and Edelman 1994; Munoz and Wurtz 1995; Munoz et al. 1991; Port et al. 2000; Soetedjo et al. 2002), which is not what we are considering here. Spatial updating across saccades is simply the associated change in activity that keeps target position in register with the external world.

[2] We continue the tradition of identifying our network by the number of its layers of units (e.g., White and Snyder 2004; Xing and Andersen 2000; Zipser and Andersen 1988). In such networks, it is the weights of the connections between or within these layers that are modified during training.

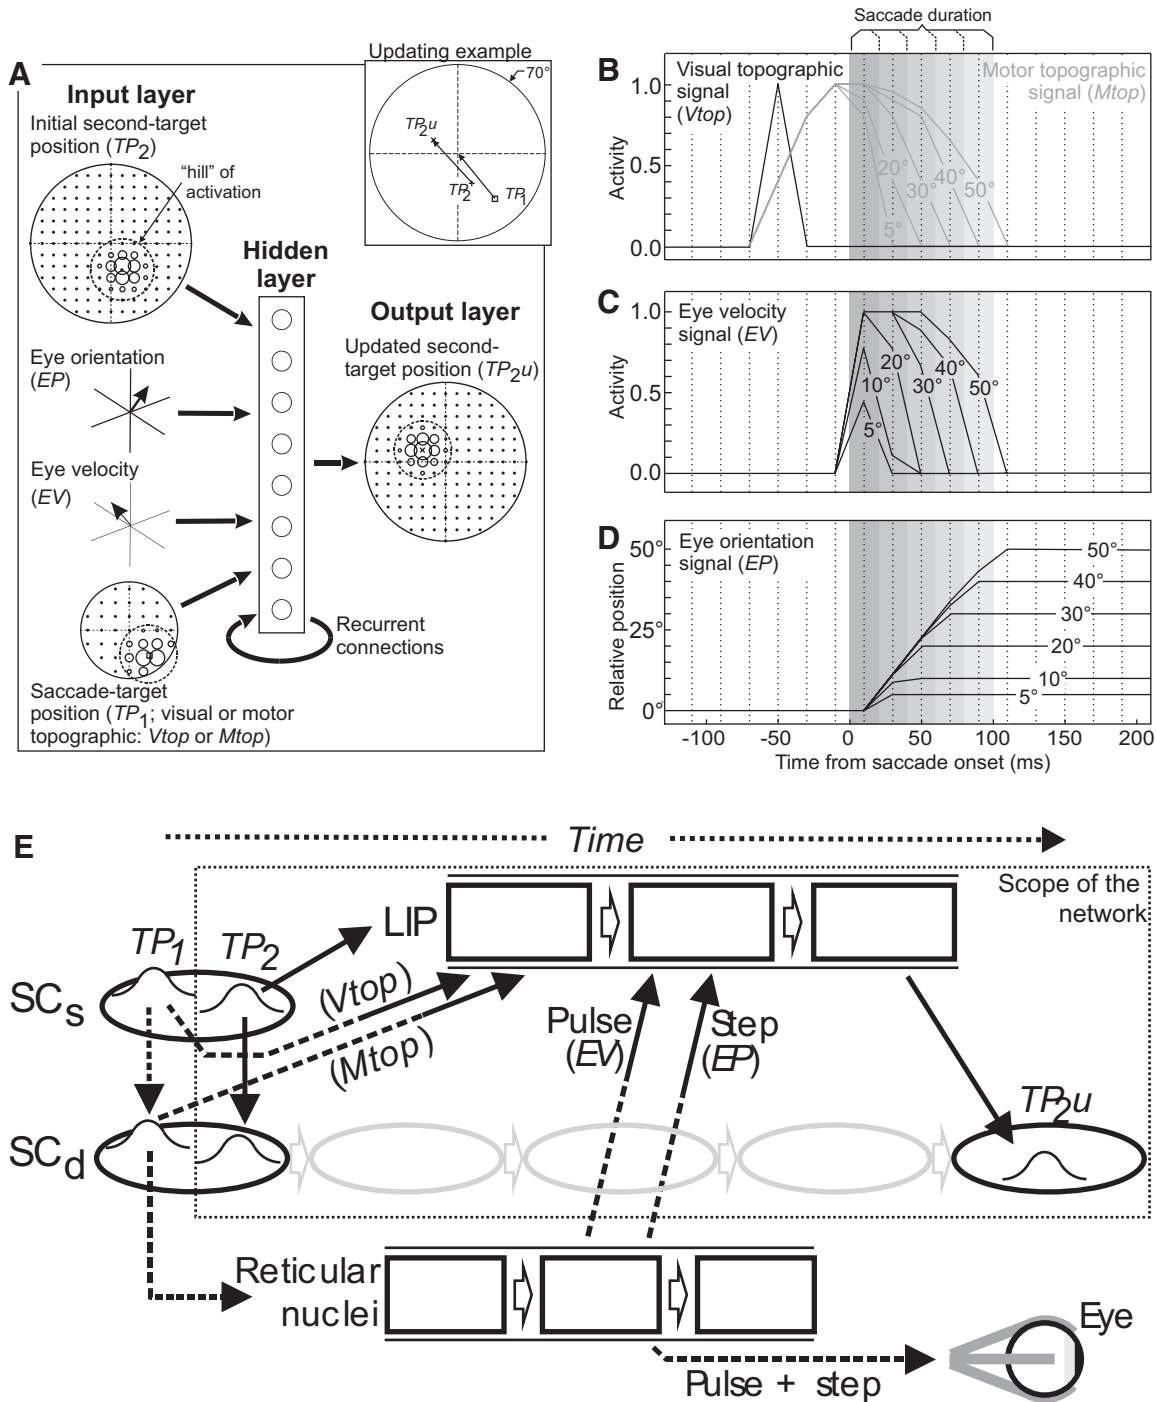[3] The online version of this article contains supplemental data.

FIG. 1.   Neural network architecture. *A*: standard network architecture showing the input and output signals for an example trial (vector diagram given in inset). Input signals are the initial 2-dimensional (2-D) 2nd-target position (TP$_2$), and the updater signals: visual or motor burst encoding the 2-D saccade-target position (TP$_1$), 3-D eye velocity (EV), and eye position (EP). The output signal is the updated (remapped) 2-D 2nd-target position (TP$_2$u). Initial and updated target positions, as well as visual or motor burst signals, are each represented by a localized hill of activity in a 2-D topographic array of units in retinal coordinates (each unit's activation for the target positions of the example trial is indicated by the diameter of the circle centered at the unit's preferred direction). Eye velocity and position are represented by rate-encoded 3-D vectors in brain stem coordinates. *B*: temporal modulation of the topographic saccade-target position signal for the visual burst (Vtop, black line) and motor burst (Mtop, gray lines for 5, 20, 30, 40, and 50° saccades) updater signals. *C*: evolution of the normalized magnitude of eye velocity signal, shown for 5, 10, 20, 30, 40, and 50° saccades. *D*: evolution of eye position magnitude (distance from primary position) for the same saccades as for eye velocity. Eye position is updated using the velocity from the previous time step. *E*: identification of the neural network with signal processing in the brain. Signal paths not modeled in the network are shown by dashed lines.

The activations of all units in the input layer were fed forward to the units of the second, hidden layer of the network. The hidden layer contained recurrent connections between all of its units. The activation of each hidden-layer unit was fed recurrently as part of the input to all hidden-layer units in the subsequent time step. Thus it was in this layer that the second-target position information was stored and updated across time. Activations of all hidden-layer units were fed forward to the units of the third, output layer of the network, which

comprised a readout of the updated second-target position (TP₂u). The updated target position in the output layer was represented in the same manner as was the initial target position in the input layer.

The initial second-target position input signal was meant to represent the burst of activity in visual neurons of either the SC or FEF that accompanies the presentation of this target. The updated second-target position signal was meant to represent the activity of motor neurons in either structure that participate in remapping (SC: Mays and Sparks 1980; Munoz and Wurtz 1995; Walker et al. 1995; FEF: Bruce and Goldberg 1985; Schall 1991; Umeno and Goldberg 1997).

TEMPORAL CODING. In our previous study (Keith and Crawford 2008), we only modeled remapping in a purely feed-forward network. In the current study, because we were concerned with the temporal evolution of signals during remapping, multiple time steps and recurrent connections were used. The temporal dimension was represented in the standard network by discrete time steps of 20-ms duration, so that the largest (50°) saccades lasted five time steps. The 20-ms time-step duration was chosen as the longest duration in which the basic evolution of the various updater signals could still be represented. We tested the possible effects of varying time-step duration on remapping in one of our network variations.

The initial second-target position signal in the input layer of the network was transient, being present only for a single time step, the first time step of each trial. The network was trained to sustain this initial second-target position in its *output* layer for the 5 time steps prior to any updater signal being presented. It was also trained to sustain the remapped activity associated with the *updated* second-target position for the five time steps *after* the updater signals ended. We used this training based on the reasonable assumption that it was only during fixation (when eye position is constant) that a coherent representation of the retinotopic target position is required. Thus we did not constrain the network's output activity in the time steps between the initial and final fixation time steps, which we termed the *updating interval*. This allowed the evolution of output activity during remapping to be an emergent property of the training for each updating model and thus potentially able to reflect the spatiotemporal properties of the different updater signals. These updating time steps comprised those during the saccade itself, three time steps prior to saccade onset during which early updater signals were presented, and one time step after saccade end before the updated target position was required.

UPDATER SIGNALS. In our previous updating model (Keith and Crawford 2008), we used a rate-encoded representation of the full motor error of the first saccade (of the double-saccade task) as the signal driving spatial updating (the updater signal). The key point of the current study was to vary the form of the updater signal (i.e., as the independent variable) to see how this affects the spatiotemporal characteristics of remapping in the hidden and output layers (i.e., the dependent variables). To do this, we selected different updater signals associated with the first saccade and provided them as inputs to the network. In theory, this updater signal could be derived at any point along the sensorimotor transformation for saccades or at more than one point. We chose three options to represent the major early, intermediate, and late stages of processing, and these three forms were used to name the three versions of our standard updating model.

Visual topographic (Vtop) model. In this model, updating was driven by a copy of a transient visual response (updater signal limited to a single time step) to the presentation of the first target in the double-saccade task. This was the target of the first saccade across which the second-target position was to be updated. The timing of the presentation of this updater signal relative to saccade onset was arbitrary. We represented it in our standard simulations as being in the third time step (50 ms) prior to saccade onset (Fig. 1*B,* black line). The spatial coding of a visual signal could take many forms, depending on where it appears in the brain. However, we were only interested in this signal in the form it takes within the saccade generator, for example,

the SC. Therefore for simplicity we used the same gaze-centered, topographic coordinates used in our output layer. It is not yet known how detailed efference copies of such signals are. We represented it in our network as being somewhat more coarse-grained than the actual topographic target-position representations of second-target position. The Vtop signal fed to the hidden layer can be thought of as a copy of the response in the superficial layers of the SC or as the transient visual burst observed in intermediate-deep layers of the SC (Mays and Sparks 1980), but it could just as well arise in cortex (e.g., FEF) (Bruce and Goldberg 1985; Schall 1991).

Motor topographic (Mtop) model. In this model, the updater signal was intended to represent a copy of the output of the SC or FEF: a saccade-related motor burst. The spatial coding for this signal was exactly the same as in the Vtop signal, but the temporal characteristics were varied with saccade amplitude based on physiological observations of such signals (Munoz and Wurtz 1995). The Mtop signal began arbitrarily three time steps (50 ms) prior to saccade onset and lasted for the duration of the saccade. The modulation of this spatially encoded signal is shown for different-sized saccades in Fig. 1*B* (gray lines). Such an efference copy has been identified passing up from the SC through the thalamus to the FEF (Sommer and Wurtz 2006).

Eye velocity (EV) model. It has also been proposed that eye velocity signals from the short-lead burst neurons of the reticular nuclei would be ideal signals to drive updating because they have the most accurate representation of the complete moment-to-moment 3-D rotation of the eye (Medendorp et al. 2003; Smith and Crawford 2001). It is not known how these signals would obtain access to all of the regions involved in remapping, but they could influence cortex through thalamic connection or via the cerebellum, which receives oculomotor signals from the brain stem and projects to the thalamus (Lynch and Tian 2005; Ohtsuk and Noda 1992). We modeled the EV signal as eye-position derivatives in head-fixed brain stem oculomotor coordinates (Crawford and Vilis 1992; Henn et al. 1989; Keith and Crawford 2008) and based their temporal profiles on the data of van Gisbergen et al. (1981). Examples of this are shown for different saccade sizes in Fig. 1*C*. We did not represent the time delay between this activity and saccade onset, which is ~12 ms (van Gisbergen et al. 1981) because the time taken for an efference copy of this signal to pass back to the cortex would have a similar delay, and the delays would thus tend to cancel. The EV signal in our network was thus nonzero only for the time steps of the saccade.

In addition to these three standard updating models, we also simulated updating models where *both* the EV and *topographic* updater (either Vtop or Mtop) signals were present.

Eye position (EP) feedback. An earlier neural network study of updating in one dimension (White and Snyder 2004) found that when such a network had access to both eye velocity and position signals, it greatly prefers the former. For this reason, we used eye velocity as the primary saccade feedback updater signal. However, our simulations of updating involved the full 3-D geometry of saccadic eye rotations because it has been shown that updating across saccades correctly accounts for the 3-D kinematics of eye orientation. This in turn requires eye orientation signals to enter the system at some point (Keith and Crawford 2008; Smith and Crawford 2001). In both the Vtop and Mtop models, the full dynamic eye position signal was required in order that the updating derived from the torsional saccade component could be carried out (because neither Vtop nor Mtop signals contain this information). In the case of the EV model, initial eye position feedback was required to relate the 3-D EV signal in head coordinates to the topographic target-position signals in eye coordinates. However, to put all three standard models on the same footing, we provided each with the full, dynamic eye position (EP) signal in addition to their principal updater signals.

The EP signal was derived by integrating the EV signal (which, again, was an orientation derivative, rather than angular velocity) and was therefore modified in each time step according to the EV in the previous time step (shown for different saccade sizes in Fig. 1*D*). The

EP signal was constant for all presaccadic time steps as it was for all postsaccadic time steps. This signal was modeled as arising from the brain stem neural integrator, which appears to use the same coordinate system as the 3-D burst generator (Crawford 1994; Crawford and Vilis 1992).

Training

The various versions of our model were trained to remap the topographically encoded activity from an initial hill of activity representing the initial second-target position. The remapped hill represented the gaze-centered activation of the same target relative to new postsaccade eye position—i.e., the updated second-target position. No constraints were placed on the output of the network during the updating interval as it was the way in which these signals evolved during updating that was a principal object of our study. The updating interval was defined as the three time steps prior to saccade onset, the saccade time steps, and one time step after the saccade.

The training set consisted of 10,000 trials, each trial composed of randomly selected initial 3-D eye position and two 2-D target positions in retinal coordinates. Saccade magnitudes were <50°. The 2-D (horizontal and vertical component) initial and final eye position were constrained to within 50° of the primary position as an approximation of the human oculomotor range. Although we only simulated saccades in head-fixed coordinates, in real life, the head also moves during gaze shifts. Thus the associated saccades have torsional eye movement components of up to $\pm15°$ (Crawford et al. 1999; Tweed et al. 1998). (These are necessary to compensate for the vestibular-driven pursuit movements that occur between gaze shifts.) The real-life updating mechanism must therefore update across saccades having nonzero torsional components (because torsion changes the gaze-centered 2-D location of any target that is not in the direction of the torsional axis). It has been shown that humans correctly update across eye/head movements with torsional components, and this places constraints on remapping that may affect the types of signals used (Medendorp et al. 2002). To capture this reality in our simulations, initial eye orientation was constrained to have a zero torsional component in keeping with Listing's law, but the final eye orientation had torsional values between $\pm15°$, randomly generated with a linear probability within this range.[4] (The full geometry and computations required for calculating the updated 2nd target position from initial and final eye position and 1st and 2nd target positions are given in the supplementary materials.)

The numbers of input- and output-layer units in our networks were constrained by the encoding scheme used for the various input and output signals, but the number of hidden-layer units was not constrained. We trained our networks with 16, 25, and 49 hidden-layer units. The analyses shown were for networks with 16 hidden-layer units, which updated target position with accuracies greater than those observed in humans (Baker et al. 2003; Herter and Guitton 1998; Medendorp et al. 2002; Schlag et al. 1990) and had the benefit of simplicity for analysis.

All of our networks were trained to perform the target-position updating using a robust variation of backpropagation-through-time (RPROP) (Riedmiller and Braun 1992). The basic approach of backpropagation-through-time training is described in the supplementary materials. Training comprised passing through the complete training set of trials (an *epoch* of training) and adjusting the weights of all network connections at the end of each epoch. This included both the weights of the connections between successive layers and the weights associated with the recurrent connections in the hidden layer. The weights of the trained network were required both to sustain the hill of activity in the output layer during fixation and remap this hill across saccades. Training was performed for 100,000 epochs for each network. To test the generalizability of our results, we tested each network on a second set of trials, generated using the same parameters as the training set but with different initial seed values and therefore comprising a distinct set.

RESULTS

Updating performance

The task required of our updating neural networks was that of sustaining in the network's output layer the initial target position (for 5 time steps) prior to the saccade and the updated target position (again for 5 time steps) after the saccade. Target position in the input- and output-layer representations was identified as the center of mass of activity in the topographic array of units of each representation. All of our updating models learned to perform target-position updating across both torsional and nontorsional (horizontal and vertical) saccade components. The root-mean-square (RMS) updating error (angular distance between actual and desired updated target positions) was measured for each model, averaged across a set of 10,000 test trials the parameters of which were similar to those of the training set with a mean saccade magnitude of 27.3°. The performance of the three standard updating model networks showed a RMS updating error of slightly over 2° (2.25, 2.20, and 2.03° for the Vtop, Mtop, and EV models, respectively), which was well within the human accuracy of such updating (Baker et al. 2003; Herter and Guitton 1998; Medendorp et al. 2002; Schlag et al. 1990).

Progression of torsional and nontorsional remapping

The progression of the activity remapping associated with target-position updating may be approximated by the movement of the center of mass of this activity during the updating. We modeled the full 3-D geometry of saccadic eye rotation in this updating task and, even though our target-position updating involved shifting only the 2-D direction of the target position (distance to the target was not considered), the torsional component of the saccade does affect updating. This *torsional updating component* comprises a rotation of target position in retinal coordinates about the fovea. The remainder of the updating (the *nontorsional updating component*) arises from the nontorsional (horizontal and vertical) saccade components. The 2-D topographic updater signal (Vtop or Mtop) provided only nontorsional updating information, while the EV and EP signals provided the full 3-D updating information. We examined how the torsional and nontorsional components of updating proceeded during spatial updating in the Vtop, Mtop, and EV standard updating models. We separated torsional and nontorsional updating components in a paired set of test trials. (Details of how this test set was constructed are given in the supplementary methods.) The fractions of the full, ideal updating components performed as of each time step for each updating model are plotted in Fig. 2, *A–C*, for large saccades (45–50°).

---

[4] Our initial training sets had a preponderance of trials with small updating components arising from torsional saccade components. This had the effect of placing the average size of updating due to the torsional saccade component at or below the accuracy of the final trained network, so that no network performed more than ~60% of this updating component. We found that by selecting trials where all sizes of updating due to the torsional saccade component were equally represented within the range produced by the variation of saccades specified in the preceding text, the full measure of updating due both to the torsional and non-torsional saccade components was learned by our networks.
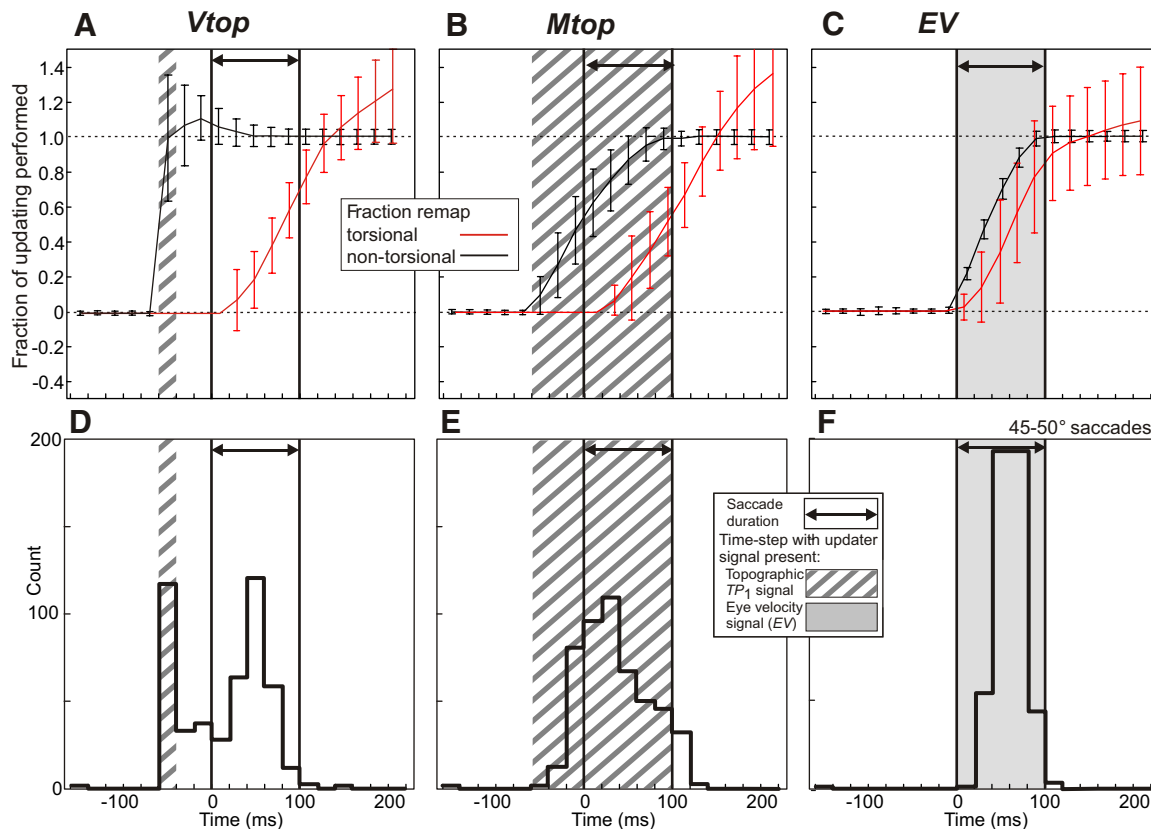
FIG. 2. Progression of remapping for large saccades. *A–C*: the fraction of updating due to the torsional and nontorsional saccade components (red and black, respectively) for the Vtop, Mtop, and EV standard updating models, plotted as of each time step (mean ± SD) for 500 large-saccade (45–50°) trials. Details of how these updating components were calculated are given in the supplementary materials. Target position was defined as the center-of-mass of activity in the output layer. Because the distributions of the spreads of fraction of updating performed were unimodal and sufficiently symmetric, SD error bars were used.*D–F*: latency spectrums are shown for the same set of trials for the standard updating models. The latency of remapped activity was defined as the 1st time step in which the activity at the ideal updated target position was >0.3 (roughly half the maximum postsaccadic activity). The activity at this position was interpolated from the activations of the 4 adjacent output-layer topographic units, weighted by the inverse of their distance from the ideal updated target position.

In the Vtop updating model (Fig. 2*A*), on average virtually all of the nontorsional updating (black) was performed in the single time step in which the Vtop updater signal was present (at −50 ms, indicated by the vertical hatched band). After this, the nontorsional updating gradually settled to the fully updated value (1.0) by the end of the saccade. The torsional updating (red) did not begin until the second time step of the saccade. This reflected the fact that dynamic EP was the only signal in the Vtop updating model that contained the torsional saccade component information that drove torsional updating. And EP only began to shift from the initial eye position value in the time following the time step in which the EV signal first became nonzero (the 1st time step of the saccade). This torsional updating continued during all successive time steps of the saccade and was not complete until the second time step after the end of the saccade.

Although torson does not settle to a completely stable position in Fig. 2, *A–C* (red), the mean torsional updating of the final time steps modeled was in all cases within a SD of the desired value of 1.0. Note that for the large saccades shown (45–50°), the torsional updating component varied from 0 to 15° and had a mean value of 4.2°. Errors in torsional updating in humans are of the order of ±2 to ±7° (Medendorp et al. 2002). The largest spread represented by the SD error bars Fig. 2, *A–C*, is ±1.3°, considerably less than human error. We arbitrarily defined complete torsional updating being achieved

as the first time step in which the SD error bars of fraction updating included the full updating value (1.0).

We confirmed that the EP signal was responsible for the entire torsional updating in the Vtop updating model by lesioning this signal (i.e., by setting it to its initial eye position value in all time steps), which eliminated all torsional updating. Furthermore, this lesioning had no effect on the nontorsional updating. Thus the Vtop signal performed all the nontorsional component of updating. Because the Vtop signal was present only in the single time step prior to saccade onset, the continuance of the residual nontorsional updating that occurred during the saccade was due to the network settling toward the fully updated target position.

In the Mtop updating model (Fig. 2*B*), the topographic updater signal began three time steps prior to saccade onset and continued until the end of the saccade (indicated by the vertical hatched band). This temporal spreading out of the topographic updater signal had the effect of spreading out the progression of nontorsional updating (black) in this updating model. Full nontorsional remapping was not complete until the end of the saccade. The torsional remapping (red), again being driven solely by the EP signal in this model, began in the second time step of the saccade and was not complete until the second time step after the end of the saccade.

We lesioned the EP signal in this model and found that the Mtop signal topographic updater signal performed all of the nontor-

sional updating, and the EP signal performed all of the torsional updating. Thus even though on average only 75% of the nontorsional updating was performed by the second saccade time step in this model when the EP signal first began to operate, the EP signal did not contribute perceptibly to the nontorsional updating and was confined to performing all of the torsional updating.

In the EV updating model (Fig. 2C), both torsional (red) and nontorsional (black) updating proceed in a roughly lock-step fashion with torsional updating lagging behind nontorsional by a little more than one time step. The nontorsional updating is completed by the first time step following the end of the saccade, whereas the torsional updating is on average completed in the following time step.

To distinguish the contributions to the updating of the EV and EP updater signals in the EV updating model network, we again lesioned the EP signal, which resulted in no detectable reduction in nontorsional remapping, and a 15% reduction of the torsional remapping. We thus confirmed in 3-D what White and Snyder (2004) found in their 1-D updating neural network: that when both EV and EP signals are present, the updating network greatly prefers the former. We also showed that when the EP signal was lesioned (in this case eliminated entirely, all input-layer units being set to their 0 values), the progression of the torsional updating after the saccade was eliminated, indicating that it was the final EP value, sustained in each postsaccade time step, that drove this residual torsional updating. That the EP signal should do this was not surprising because the EV signal is zero after the saccade (the same value it had prior to saccade onset), while the postsaccade EP signal is different from its presaccade value.

That the topographic updater signal in these models should perform only the nontorsional updating was not surprising because it did not contain any information about the torsional saccade component. What was interesting, however, was that for the models containing the topographic updater signal (Vtop and Mtop), there was more or less complete separation with the topographic updater signal performing only the nontorsional updating and the EP signal performing only the torsional updating. In the EV updating model, on the other hand, the EV signal performed both all of the nontorsional and most of the torsional updating. Clearly, while the EP signal is of secondary importance in the EV model, it *is* important in the topographic updater models because it is the only updater signal that contained torsional saccade information.

Because of the onset of the topographic updater signal *prior* to saccade onset, the nontorsional updating started before the saccade in both Vtop and Mtop models. Thus in these models, nontorsional updating preceded torsional updating, although this distinction is clearer for the Vtop model. In the EV model, on the other hand, torsional and nontorsional updating proceeded roughly at the same time. Thus we expect curved updating paths for the Vtop and Mtop models for trials in which the torsional updating was significant and in a direction different from that of the nontorsional updating. We did in fact see this for some trials but noted that this was only the path as represented by the center of mass of activity. In the Vtop and Mtop models, the activity during updating was not a coherent hill of activity. This made the detection of this curvature, and thus the differentiation between Vtop or Mtop and EV updating models by comparing the latencies of activity of different output-layer topographic units (such as at the nontorsional-only vs. fully updated target positions), statistically impossible.

The detection of presaccadic remapped activity (interpolated from the activity at units adjacent to the updated target position in our models) was a principal finding of neurophysiological remapping studies (Duhamel et al. 1992; Umeno and Goldberg 1997; Walker et al. 1995). In Fig. 2, *D–F,* we plot the latency of remapped activity (i.e., activity at the updated target position) for 500 large-saccade (45–50°) trials for all three standard updating models. We see that while all models produce latency spectrums that span the saccade and immediate postsaccade intervals, only the Vtop and Mtop models (Fig. 2, *D* and *E*) show presaccadic remapped activity (i.e., activity at the updated target position prior to saccade onset). The EV model fails to produce such activity (Fig. 2F). This fact alone suggested that remapping in the brain cannot be driven solely by the EV signal.

What is also of interest is the bimodal remapped-activity latency spectrum for the Vtop model (Fig. 2D). Because the entire nontorsional updating has on average been performed as of the first time step of this spectrum (and torsional updating is generally much smaller than nontorsional updating), this bimodal property cannot be explained solely by the movement of the center of mass of activity. Clearly there is something else going on during the updating period. This will be discussed in a later section, but first we will analyze the behavior of the hidden-layer unit activities during this updating.

Emergent structure of the hidden layer

We have seen that the updater signals used affect how each model solves the updating problem in terms of the progression of torsional versus nontorsional updating and remapped activity latencies. To understand how these differences came about, we first analyzed how the hidden layer of each network was organized and how it functioned. Because the hidden-layer units were not a priori organized topographically, our spatial analysis of hidden-layer activity during updating used the spatial properties of individual hidden-layer units—i.e., the remapping of their receptive fields. We also focused on functional properties of the hidden layer such as receptive and motor fields rather than statistical analyses of connection weights because the former are more appropriate for physiologists.

TARGET POSITION RETENTION BETWEEN SACCADES. Our recurrent networks were trained to sustain a target position in their output layer in those time steps in which no updater signal was present—i.e., initial second-target position before and updated second-target position after the updating interval. Because the initial second-target position was presented to the network for one time step only (the 1st time step of each trial), this position information was retained by the population activity of the network's hidden-layer units by means of the recurrent connections. These recurrent connections fed the activations of each hidden-layer unit in a given time step back as inputs to all hidden-layer units in the following time step.

During memory intervals without saccades, the networks relaxed into steady states of hidden-layer unit (HLU) activation (which had the property of generating a hill of activity in the output layer centered at the second-target position being re-

tained). The mechanism by which this steady state was maintained was basically that described by Xing and Andersen (2000): the recurrent weights within the hidden layer were positive between HLUs having similarly oriented receptive fields and negative between units with very differently oriented receptive fields. Details of these results are shown in the supplementary materials. Here we focus on intrasaccadic remapping.

HOW THE UPDATER SIGNALS DROVE REMAPPING. The network was trained to sustain a given target position by means of a hidden-layer state that was maintained through a set of recurrent connections with learned weights. The addition of the updater signals input to the hidden layer acted to perturb this state so as to produce a change, remapping activity so as to update this target position. Because no target position was specified in the output layer during the updating interval, the properties remapping were an emergent property of the network in both hidden and output layers.

We examined how the receptive field of HLUs evolved during updating and found that, for the most part, these receptive fields retained their basic structure, but were shifted—i.e., remapped across successive time steps. We quantified this remapping as the shift of receptive field features (such as an iso-activation curve, defined as all 2nd target positions for which the HLU's activity was a specific value) relative to the time step immediately prior to the start of updating (70 ms before saccade onset). Figure 3 shows the horizontal remapping of all HLUs in each standard updating model for an example trial involving updating across a 30° rightward saccade. All three standard updating models showed full remapping (30°) in all HLUs by the second time step after the saccade (90 ms after saccade onset). The models differ, however, in the progression of remapping of their HLUs during the updating interval (-50 to $+70$ ms).

In the EV updating model (Fig. 3C), the progression of remapping for all HLUs was highly regular during the updating, with values more or less limited to the updating range (0–30°). These curves did vary in terms of their *timing*: some units completed their remapping by the middle of the saccade while others did so only in the second time step after saccade end.

In the Mtop updating model (Fig. 3B), the spread of remapping values across all HLUs during updating was not limited to the updating range. Some units showed transient remapping in the direction opposite to the updating ($<0°$), whereas others showed a transient remapping that overshot this updating ($>30°$). The spread of remapping within each time step increased during the period immediately prior to saccade onset and diminished toward the end of the saccade. In fact, the size of the remapping spread among the HLUs roughly matched the modulation of the Mtop updater signal (see Fig. 1B, profile for 30° saccade).

In the Vtop updating model (Fig. 3A), the spread of remapping during the updating interval was even greater than for the Mtop model. In the time step in which the transient updater signal was present (-50 ms, vertical hatched band), the range of HLU remapping covered the entire measurable range ($\pm100°$). After this the spread of remapping rapidly diminished, becoming close to zero by the first time step of the saccade ($+10$ ms).
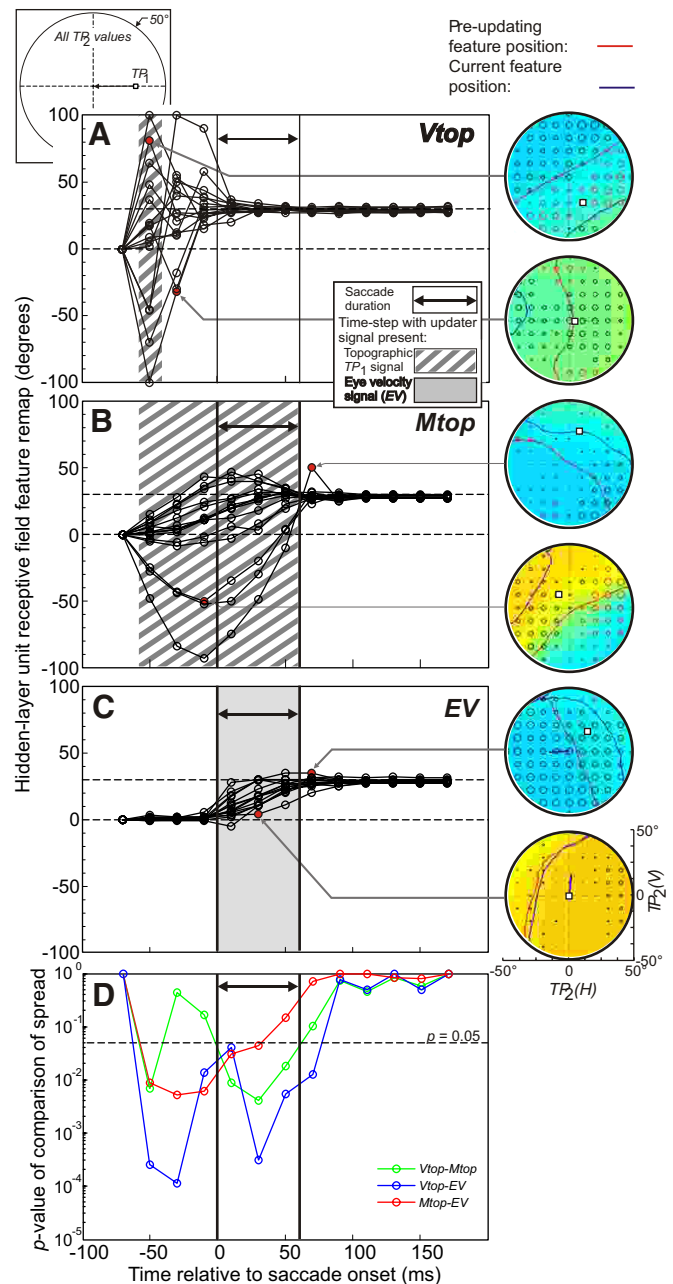

FIG. 3. The remapping of hidden-layer unit (HLU) receptive field populations across a 30° saccade. The horizontal remapping of the receptive field of all HLUs is plotted during updating across a 30° rightward saccade. Initial eye position was primary position. Remapping was measured as the shift of a receptive field feature relative to its position in the time step immediately prior to the updating period (70 ms before saccade onset). The feature used was a curve of constant activation (all initial 2nd-target positions for which the HLU's activation had a certain value). Remapping shifts are given for each subsequent time step for the Vtop (A), Mtop (B), and EV (C) updating models. The limits of the remapping were $\pm100°$. Also shown are 2 example HLU receptive fields at a single time step for each model, where the position of the feature in that time step is shown as a blue curve and its position in the preupdating time step is shown as a red curve.

Thus our three standard updating models predicted very different patterns of population remapping in their HLUs: for the EV model, there was little spread, for the Mtop model, there was a great deal of spread that lasted from before saccade onset until the end of the saccade, and for the Vtop model,

there was an even larger spread that all but disappeared by saccade onset. This means that the signal used to drive updating may be inferred from the pattern of transient remapping observed among a population of neurons in a given brain area.

GAIN FIELD MECHANISM OF HIDDEN-LAYER REMAPPING. The way in which the updater signals affect the HLU activations during updating may be represented using gain fields: modulations of receptive fields first observed for neurons by Andersen et al. (1985) as a function of eye position but also proposed as the mechanism for updating (Cassanello and Ferrera 2007; Keith and Crawford 2008; White and Snyder 2004). Neural network simulations have the advantage of allowing us to examine the updater-signal input that is *responsible* for the gain-field activity modulation in a given HLU rather than simply the resulting modulation that comprises the gain field.[5] This allowed us to observe the full spectrum of the updater signal's input to each HLU as a function of saccade-target position (TP_1).

We show examples of *updater input fields* for two HLUs in each standard updating model in Fig. 4. These were generated by plotting the total updater-signal input to a given HLU (the sum of the activations encoding the updater signal in the input layer, each activation multiplied by the associated weight of the connection to the HLU), for all different saccade-target position (TP_1) values within an eccentricity of 50°. It was immediately apparent that the updater input fields for the Vtop and Mtop models (Fig. 4, *A* and *B*) were highly irregular, but those for the EV model (*C*) were highly regular, for all HLUs in these models.

The regularity of the EV model updater input field is reflected in the fact that this field for each HLU may be completely described by the *sensitivity vector* of that unit (Keith and Crawford 2008) and the EV signal modulation curve of Fig. 1*C*. The definition of this sensitivity vector and how it describes the EV updater input field are given in the supplementary materials. The sensitivity vectors are shown in Fig. 4*C* as blue vectors originating at the center of each field. These vectors define a line of perfect symmetry in the field. The regularity of these updater input fields may be shown in a simpler fashion by taking a slice through them along the horizontal meridian (dashed or solid horizontal red line across each field). The variation of the updater-signal input value with horizontal TP_1 is shown for the two example HLUs in Fig. 4*C*, *bottom*. For each HLU, these values vary monotonically and symmetrically about zero and saturate at just 20°.[6] The important feature of these EV updater input fields is that they were highly regular for all HLUs.

The updater input fields of the Vtop and Mtop updating models, on the other hand, were highly irregular (Fig. 4, *A* and *B*). The only basic spatial parameter that could be defined for each updater input field was its center of mass (indicated by the white square in each field). This center of mass indicated something of the magnitude and direction of TP_1 that was the weighted target position mean that the updater signal favored. Again horizontal slices through the example updater input fields are shown in Fig. 4, *A* and *B, bottom*. The variation of

---

[5] The gain-field modulation of a neuron is limited by the threshold and saturation properties of the activity. This property is reproduced in our model's hidden- and output-layer units by the effects of the sigmoidal transfer function that squashes the total input to a given unit into a limited range of activation—zero to one—compressing all large inhibitory or excitatory total inputs into values near these limits.

[6] This saturation reflects the fact that the updater input fields shown in Fig. 4*C* are those associated with the first time step of a saccade. For saccade magnitudes of >20° the magnitude of the EV signal in the first time step is at saturation (Fig. 1*C*). The EV signal for larger saccades is represented by increasing the modulation value of the EV signal in subsequent saccade times steps.
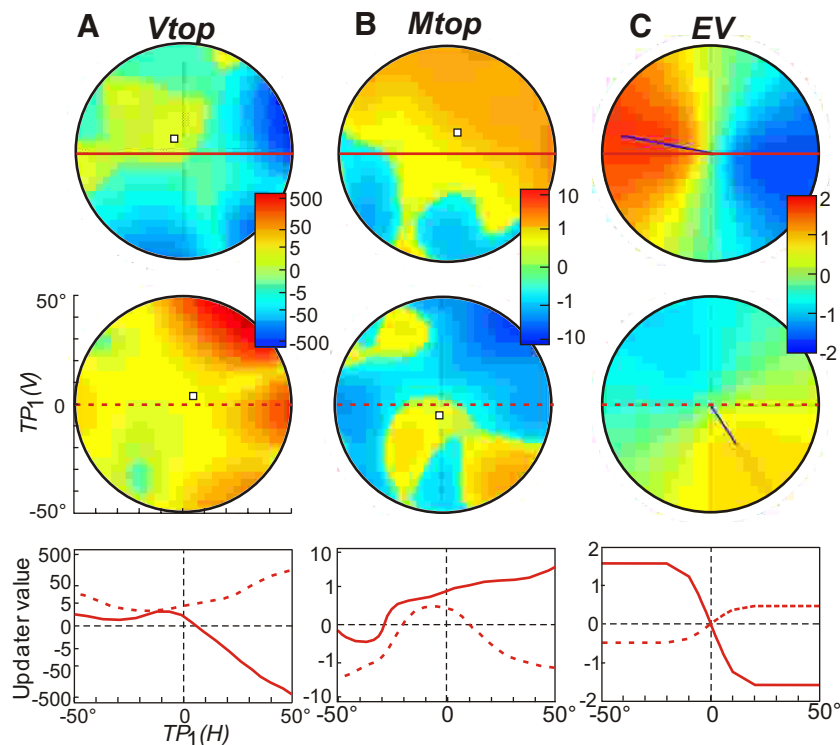

FIG. 4. Examples of HLU updater input fields. The total updater input to 2 example HLUs are shown for the Vtop (*A*), Mtop (*B*), and EV (*C*) standard updating models, plotted as a function of 2-D saccade-target position, TP_1. Also shown are the centers of mass of each field for the Vtop and Mtop models, and the 2-D sensitivity vector (blue vector) for the EV model. *Bottom row*: slices of the updater input fields along the horizontal meridian are plotted for the example HLUs.

updater-signal input with horizontal saccade-target position can be seen to be nonmonotonic and highly irregular.

But how does the regularity (vs. lack of regularity) of the updater input fields affect the remapping of HLUs in a given updating model? The monotonic and regular horizontal slices of the EV model updater input fields meant that as the horizontal first-target position is gradually increased or decreased, the updater-signal input either increased or decreased in a regular fashion (up to saturation). This is consistent with systematic, regular remapping because the lateral shift in a given HLU's receptive field was proportional to the updating signal applied to it (Keith and Crawford 2008). On the other hand, the irregular horizontal slices of the Vtop and Mtop model updater input fields meant that, as horizontal first-target position is gradually increased or decreased, the updater input field value changed in an unexpected fashion, in direction, and even sign. This irregularity did *not* favor regular patterns of remapping for the HLUs. The different variation of each HLU updater input field with saccade-target position suggests instead a dispersed pattern of remapping across the HLUs during the updating period.

FORM OF REMAPPING IN HLUS. The receptive and motor fields that the hidden-layer units developed in all three standard updating models for some HLUs bore similarities to those developed in the HLUs of our earlier feed-forward network (Keith and Crawford 2008). For example, the relationship between the receptive and motor field of a HLU, where the "on" regions of the unit's receptive field (the regions of significant activation) tended to correspond to the excitatory regions of the same unit's motor field (indicated by small red circles), was retained for some HLUs. This is seen in the example HLUs of Fig. 5, *A* and *C,* in the time step just prior to the onset of remapping (*2nd row* from the *top*) but not for Fig. 5*B*. This relationship was observed in the earlier feed-forward updating network due to the way in which HLUs combined their contributions to produce a hill of activation in the output layer (described in detail in Keith and Crawford). The fact that this general pattern was retained only in some HLUs in the updating models of the current study meant that something else was going on.

One feature of the evolution of the HLU receptive fields was that during the first five time steps, the receptive field gradually simplified, changing from a complex structure in the first time step (at -150 ms) to one that was much simpler by the fifth time step (at -70 ms). Examples of this are seen in the examples of Fig. 5 (compare the *top* and *2nd rows*). This change occurred as the network settled into a steady state while it continued to generate the initial second target position in its output layer. The reduction in complexity reflected the decrease in the number of units involved, from the 149 input-layer units representing the initial second target position in the input layer of the network in the first time step to the 16 HLUs used to sustain (and update) target position in the hidden layer in subsequent time steps. Networks having more HLUs had somewhat more complex receptive fields in their fifth time step.

In terms of the HLU receptive field remapping, the earlier feed-forward network (Keith and Crawford 2008) had the updater signal acting only directly, and therefore *separately*, on each HLU by means of the feed-forward connections from the updater signal in the input layer to each HLU. In that network, all HLUs showed the same remapping—that associated with

the full updating of the second target position (because the network performed updating in a single time step).

In the recurrent networks of the present study, the ultimate origin of all remapping was again the updater signal in the input layer. However, the recurrent connections between all HLUs created a communication between all HLUs across successive time steps during the updating. This meant that the remapping seen in individual units of all but the first time step of the updating depended on the changes in activation in the entire *population* of HLUs. And this had an effect on various properties of these units.

An example of this is shown in the evolution of remapping of the example HLU of the EV updating network in Fig. 5*C*. The EV updater signal was present only during the time steps of the saccade (indicated by the gray background of time steps at 10, 30, and 50 ms). The 2-D sensitivity vector is superimposed on the receptive field for all time steps as a blue vector pointing from the center of the field. As may be seen for this HLU, this vector points down and to the left. Because the saccade in this example was 30° to the right, and thus purely horizontal, the EV updater signal during this saccade was negative (depending as it does on the dot product of the sensitivity and saccade vectors, described in detail in the supplementary materials). A negative updater signal will drive a sigmoidal receptive field such as is seen in Fig. 5*C* in the direction *opposite* that of its downward gradient. This is what is observed in the first time step having a nonzero EV (10 ms). Here the current feature position (blue curve) has been shifted to the left relative to its position in the preremapping time step position (red curve). That is, in the direction opposite to that of the saccade vector (black arrow). This "backward" remapping, however, occurs only for this single time step. In subsequent time steps, the receptive field feature (blue curve) shifts increasingly to the right—i.e., in the "forward" direction. Because the effect of the EV updater signal is *defined* by the sensitivity vector for the HLU in each time step, this effect itself must always be in the leftward, "backward" direction for this HLU. This means that the "forward" remapping in this unit was driven not by the direct effect of the EV signal to this unit, but indirectly through the changing activations of the other HLUs, acting through their recurrent connections to this unit. Thus the remapping of individual HLUs is driven by the changing activation of the full population of these units.

An important feature of this population-driven remapping is that one limitation of the feed-forward network does not apply to our recurrent updating neural networks, namely that only sigmoidal or open receptive fields may participate in the remapping in the hidden layer. And this means that fully remapping closed receptive fields may be present in such networks. This indeed turned out to be the case. For all updating models, some HLUs developed receptive fields that were nonsigmoidal, in some cases, fully closed. This may be seen in the example HLUs shown for the Vtop and Mtop updating models in Fig. 5, *A* and *B*. For these HLUs, both the leading and trailing edges are remapped in the same desired direction.[7]

---

[7] A HLU having a closed receptive field, driven only by direct updater signals such as occurs in feed-forward networks, would show remapping in opposite directions on its opposite edges because the gradients on these edges would be opposite.
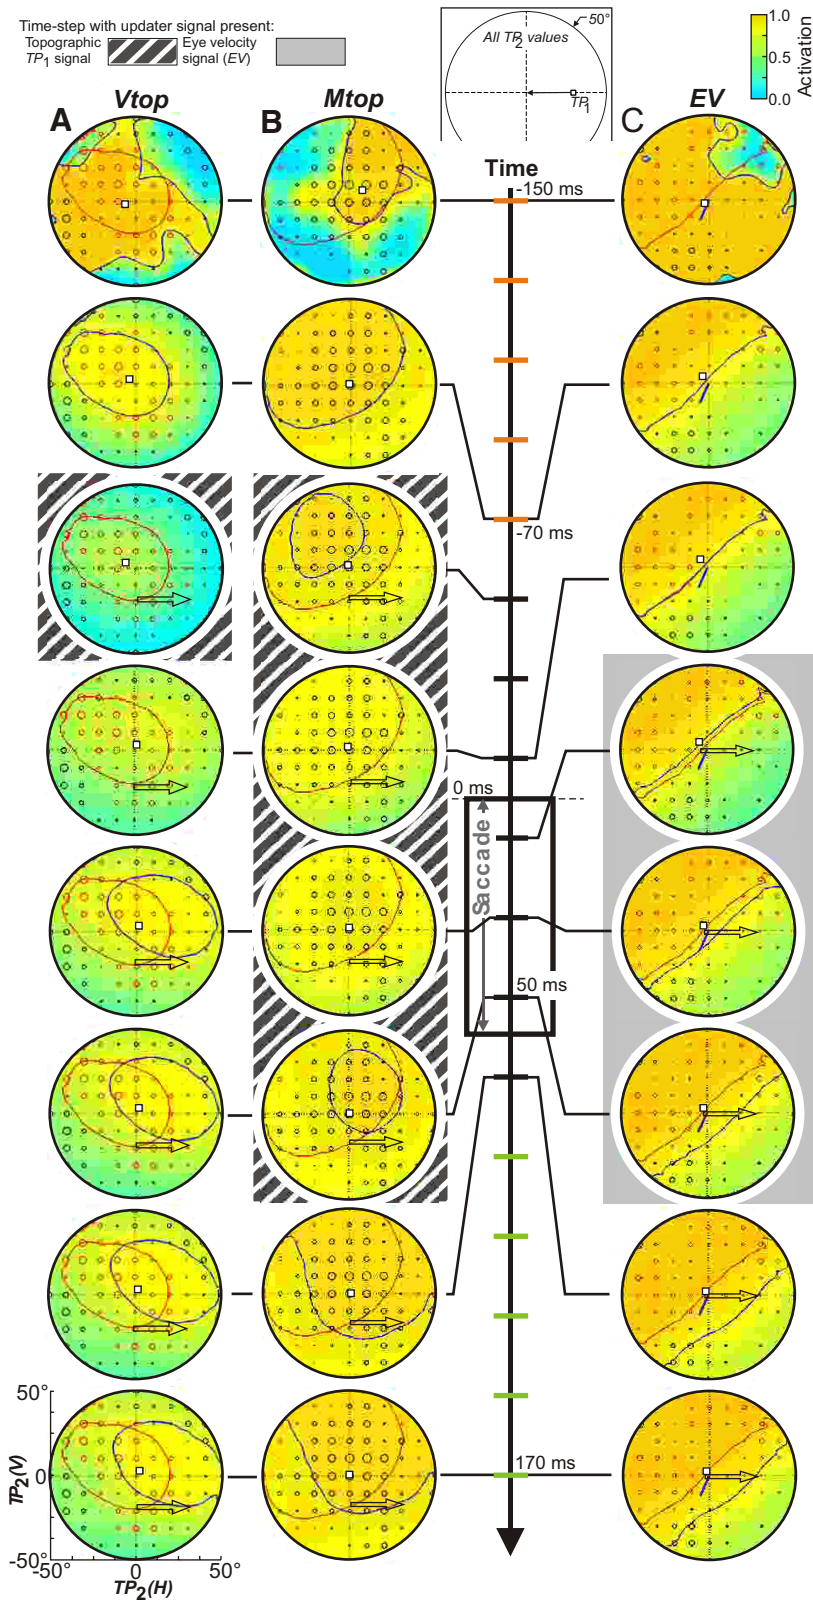
FIG. 5. Evolution of example HLU receptive fields during updating. The receptive fields of example HLUs are shown in selected time steps across updating produced by a 30° rightward saccade for the Vtop (*A*), Mtop (*B*), and EV (*C*) standard updating models. Time progresses from the 1st time step of the trial (at 150 ms prior to saccade onset) at the top to the last time step of the trial (at 170 ms after saccade onset) at the bottom. Each receptive field is a plot of the unit's activation in that time step as a function of initial 2-D target position, TP_2. Superimposed on each receptive field is the center of mass of the field (white squares). Also shown are the motor fields of the HLUs, defined as the weights connecting the HLU to each output-layer topographic unit and plotted as small circles at the preferred direction of each topographic unit, red = positive and black = negative, with the diameter of the circle proportional to the magnitude of the weight. For the EV updating model, the 2-D sensitivity vector is shown as a blue vector. Also plotted in each time step is a red curve of constant activation that was situated near the center of the field in the pre-updating time step (time step 5, at −70 ms), and a blue curve showing the same activation value in the current time step. The 30° rightward expected remapping vector is shown as an arrow.

As we have seen, our three standard updating models, which used different efference copy signals to drive the required target-position updating across saccades, developed hidden layer properties that could be distinguished experimentally. While individual HLUs could develop open or closed receptive fields, the features of each receptive field showed remapping in the postsaccade period that was tightly consistent across all HLUs and all models, equal to the saccade movement vector. During the updating period, however, when the Vtop signal was used to drive updating, HLUs showed a wide range of

transient remapping values only prior to saccade onset, a range that extended far beyond the range of the initial-to-final remapping positions. When the Mtop signal was used to drive updating, a similarly large transient range of remapping developed across its HLUs, but while this spread again started prior to saccade onset, it continued for the duration of the saccade as well. When the EV signal was used to drive updating, on the other hand, all of the HLUs remapped in a fairly coherent fashion, approximately limited to the initial-to-final remapping range. Thus the signal used to drive target-position updating should be reflected in the pattern of remapping observed across neuron populations within a given brain region.

Emergent properties of the output layer

The output-layer activity prior to and after the updating interval, comprising the first and last five time steps of each

trial, were constrained to represent initial and updated target positions, respectively, as coherent hills of activity centered on these target positions. The unspecified, and therefore emergent, properties of the output layer activity were thus limited to the time steps during the updating period, corresponding to just before, during, and immediately after the saccade across which the updating was performed.

EVOLUTION OF OUTPUT-LAYER ACTIVITY.　Figure 6 shows the evolution of activity of the entire output layer for an example 50° saccade trial for the three standard updating models. Only single pre- and postremapping time steps are shown (the *top* and *bottom panels* for each model). All three models show the required well-defined hill of activity centered at the initial second target position (open circle) in the preremapping time step, and at the updated second target position (closed circle) in
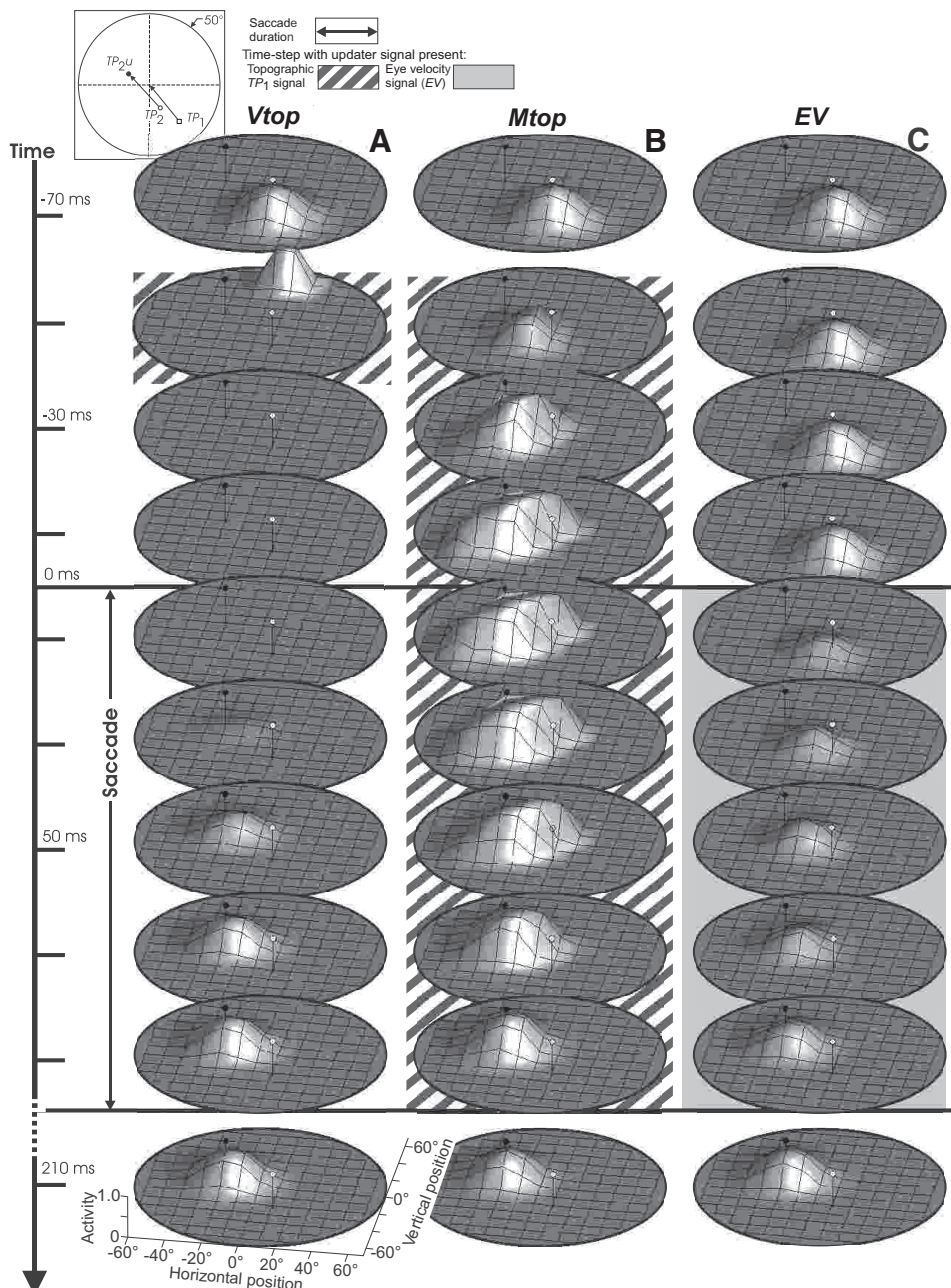

FIG. 6. Output-layer activity remapping for an example trial in the standard updating models. The evolution of output-layer activity is plotted for an example large-saccade updating trial (shown in *inset*) for the Vtop (*A*), the Mtop (*B*), and EV (*C*) standard updating models. The activations of all output-layer units are plotted in each time step as a function of the 2-D preferred directions of these units. Initial (○), TP$_2$, and updated 2nd target positions (●), TP$_2$u, are also shown.

the postremapping time step. The output-layer activities differ significantly, however, in the intervening time steps of the updating interval.

The Vtop updating model (Fig. 6A) shows a transient agglomeration of activity that is far from either initial (open circle) or updated (closed circle) target positions, during the time step in which the Vtop updater signal is present (at −50 ms, indicated by hatched background). In subsequent time steps, there is more or less complete suppression of all output-layer activity until the second time step of the saccade (at 30 ms), when a hill of activity starts to appear close to the updated target position (closed circle). In the remaining time steps of the saccade, this hill of activity grows in amplitude. This is an example of the remapping of output-layer activity in the Vtop updating model: a suppressed jumping hill preceded by a transient phase of highly sporadic activity.

The Mtop updating model (Fig. 6B) in the first time step where the Mtop updater signal is present (these time steps indicated by hatched background) shows a small jump partway toward the updated target position (closed circle). In subsequent time steps, the hill spreads in the direction orthogonal to the remap vector (from open to closed circle) while the center of mass of this activity continues to move steadily toward the updated target position (closed circle). This is an example of the remapping of output-layer activity in the Mtop updating model: a moving, somewhat dispersed hill of activity.

The EV updating model (Fig. 6C) shows the hill of activity being coherently sustained at the initial target position during the time steps immediately preceding the saccade (during which the EV updater signal is not present). This is followed by a moderate suppression of the activity at the start of the saccade that is gradually removed as the saccade progresses. During these saccade time steps, a coherent hill is maintained, gradually moving toward the updated target position (closed circle). This is an example of the remapping of output-layer activity in the EV updating model: a coherent moving hill of activity.

EVOLUTION OF OUTPUT-LAYER RECEPTIVE FIELDS. Complementary to the representation of remapping in terms of the entire population of the output layer is that of individual output-layer unit (OLU) receptive fields. This receptive field remapping is shown for an example OLU for each standard updating model in Fig. 7 across 30° leftward, horizontal saccades for all initial second-target positions. In the pre- and postupdating interval time steps shown, the OLU receptive field was a Gaussian shape as it was defined to be. The shape of the receptive field during remapping was widely divergent for the three updating models.

In the Vtop model (Fig. 7A), the OLU's receptive field spread greatly in the time step where the Vtop updater signal is present (indicated by hatched background). In subsequent time steps, the receptive field was almost totally suppressed but then gradually re-emerges, centered at a position nearly 30° to the left of the unit's initial receptive field preferred direction.

In the Mtop model, the unit's receptive field gradually spread in the presaccadic time steps, and then gradually re-coalesced in the saccade time steps to the 30°-to-the-left shifted position. In the EV model, a coherent Gaussian

receptive field was maintained in all time steps with the position gradually shifting to the left.

The movements of these receptive fields are relevant to neurophysiologists, who, by measuring the activity of individual neurons across different translated trials, can sample different locations within the neuron's dynamic receptive field.

DESCRIBING THE ACTIVITY DISTRIBUTION. From example trials of Fig. 6, we saw that the output-layer population activity cannot be fully described simply in terms of the position of its center of mass of activity. We therefore developed a set of four descriptors that described the basic features of the evolution of this population activity during the remapping period. These descriptors are shown in Fig. 8 (mean ± SD) for a set of 500 large-saccade (45–50°) trials, the largest saccades being used because the effects observed were most marked for these trials.

The fraction-of-remapping-performed descriptor, already seen in Fig. 2, A–C, was now used to represent the fraction of total remapping performed as of each time step. This descriptor was defined as the vector from the initial target position to the current activity center of mass, projected onto the ideal remapping vector (from initial to ideal updated target position), normalized by dividing by the size of the latter vector. This *fraction remap* descriptor, shown in Fig. 8A (black) for all models, had the same evolution as the nontorsional updating of Fig. 2 because the torsional updating was generally much smaller nontorsional updating for all trials.

The maximum activation was the largest activation value of all output-layer units in each time step. This *max activation* descriptor, shown in Fig. 8A (red) for all models, showed a strong suppression of activation in the Vtop model in the time steps immediately following that in which the Vtop updater signal was present (indicated by hatched background), while both Mtop and EV models showed only a slight suppression at the onset of their respective updater signals.

The lateral shift of the center of mass was the perpendicular distance of the center of mass of activity to the straight-line remapping vector (from initial to ideal updated target positions—i.e., from the center-of-mass position of the initial hill of activity and the ideal remapped hill of activity). This *lateral shift* descriptor was shown in Fig. 8B (black) for all models. In the Vtop model, it showed a very large increase (on average of 20°) in the time step where the Vtop updater signal was present (hatched background), representing the highly sporadic position of the activity in this time step. In subsequent time steps, this lateral shift decreased as the center-of-mass position settled toward the updated target position. In the Mtop model, the lateral shift increased during the presaccadic time steps where the updater signal was present (hatched background) and then decreased toward the end of the saccade. This evolution roughly followed that of the Mtop updater signal modulation profile shown in Fig. 1B for 50° saccades. In the EV model, there was only a slight increase in lateral shift during the saccade when the EV updater signal was present (gray band). The residual increase in the postsaccade lateral shift seen in all models arose from the on-average several-degree updating error shown by these models.
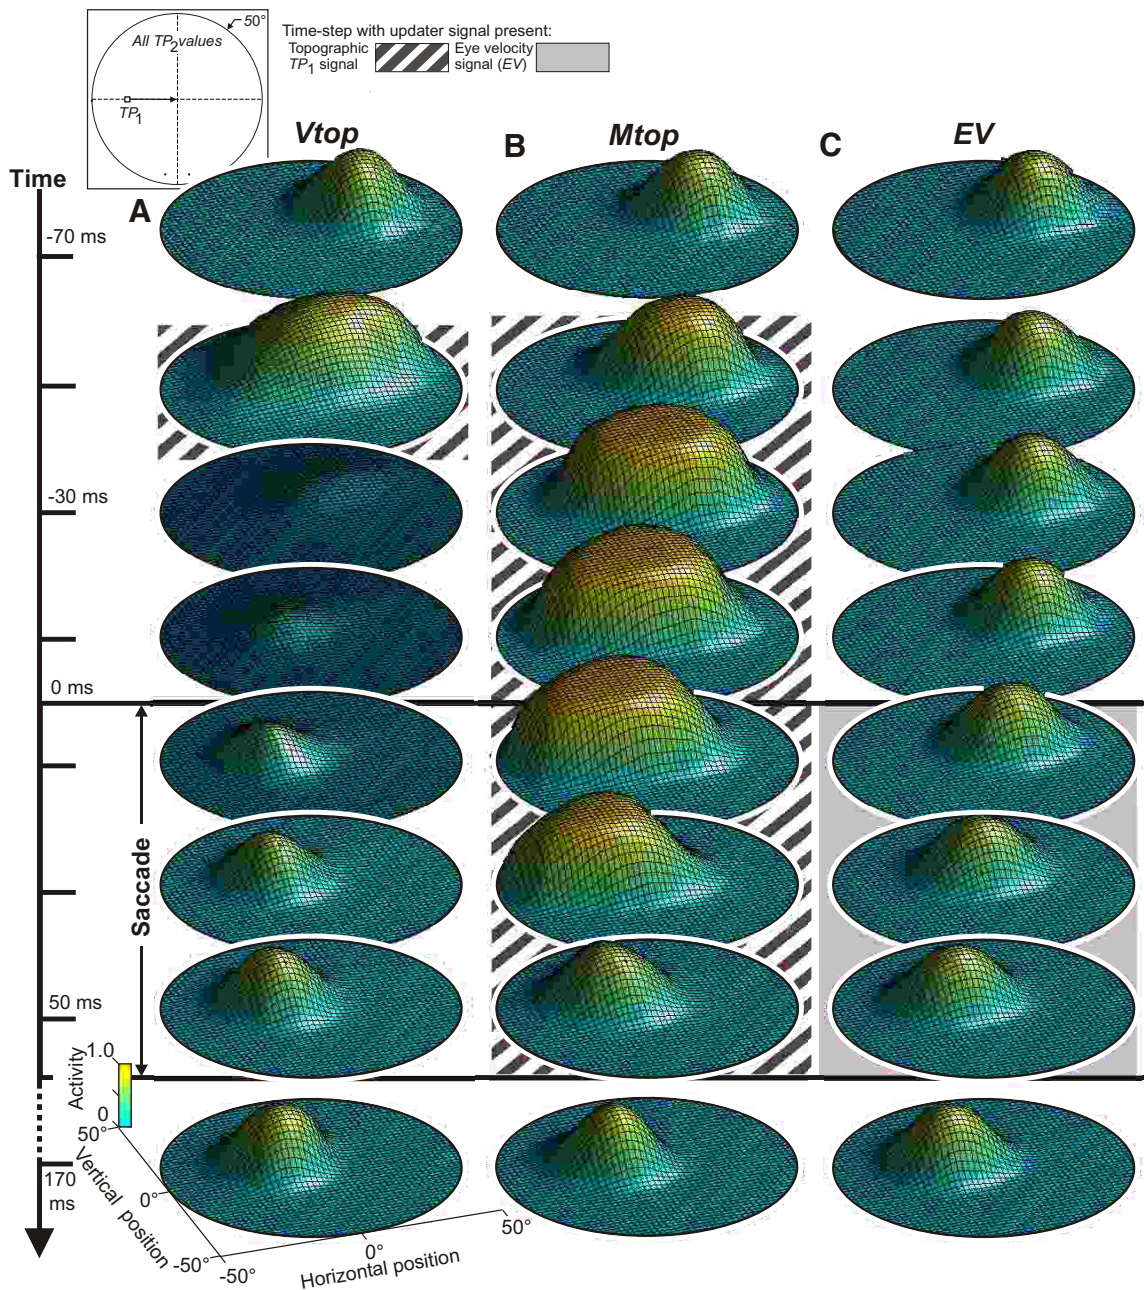
FIG. 7. Evolutions of example output-layer unit receptive fields during updating in the standard updating models. The evolution of receptive fields of example output-layer units is plotted for updating across 30° leftward saccades for the Vtop (*A*), the Mtop (*B*), and EV (*C*) standard updating models. These fields are plots of the activation in each time step of the example output-layer unit plotted as a function of the initial 2nd target position, which was varied across all values within an eccentricity of 50°.

The spread of the activity about the center of mass was represented by the root-mean-square distance of all output-layer activity from the activity center of mass, where unit distances were weighted by activation.[8] This *spread* descriptor is shown in Fig. 8*B* (red) for all models. The baseline spread, indicated by the horizontal dashed line, corresponded to an ideal hill of activity defined by the units' receptive fields. In the Vtop model, the spread increased,

---

[8] We did not distinguish spread that was parallel or perpendicular to the direct remap vector for this spread descriptor and only made this distinction for the center of mass because movement of the center-of-mass position perpendicular to the direct remap vector is not related directly to the updating.

and became highly variable across trials, in the time step where the Vtop updater signal was present (hatched background). In subsequent time steps, the spread decreased both in terms of average and in variation toward the baseline level. In the Mtop updater model, the spread increase was more moderate but lasted for the duration of the Mtop updater signal (hatched background), again roughly following the modulation profile of the Mtop updater signal. In the EV model, there was only a slight increase in spread during the saccade.

To evaluate whether the changes in max activation, lateral shift or spread descriptors during remapping were robust
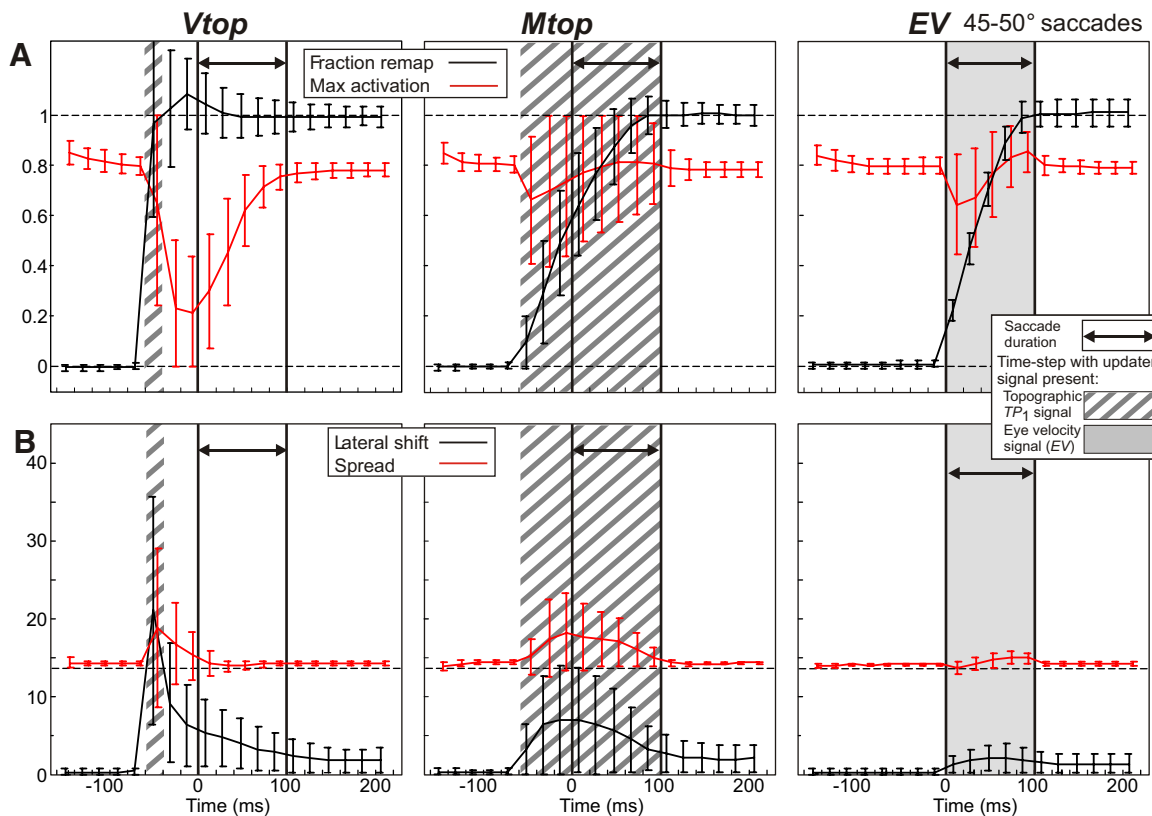
FIG. 8. Evolution of four descriptors of output-layer activity across updating in the standard updating models. The means ±SD for 4 descriptors of output-layer activity from a set of 500 large-saccade (45–50°) updating trials are plotted for the Vtop, Mtop, and EV standard updating models. SDs were used because the distributions of spread of descriptor values were unimodal and sufficiently symmetric. The single exception to this symmetry was for the activation descriptor when mean values approached either limit (0 or 1), in which conditions the error bars were modified accordingly. This was done in all figures showing activation spreads. A: the fraction of updating performed (black) as of each time step is defined as the normalized projection of current remapping onto the ideal remapping vector. The maximum activation (red) is the largest activation of all the output-layer units in the current time step. B: lateral shift is the distance of the center of mass of output-layer activity (black) in degrees from the direct remapping vector from initial to updated target position. Spread of activity (red) in degrees is the root-mean-square distance of all output-layer unit activity from the center of mass of this activity.

and significant, we compared their values in the last preupdating time step (the last postupdating time step for the lateral shift descriptor)[9] with each of the time steps of the updating interval. Because P values generally become more significant for larger sample sizes, we chose quite a small sample size—in fact that of the data reported by Sommer and Wurtz (2006) for remapping midpoint activation: $n = 13$. With such a small sample size, however, the P value of the t-test will vary significantly, depending on the updating vectors of the trials in the sample. We therefore tested multiple distinct samples of 13 trials each, and constructed a 95% confidence interval (CI) for the P values from all samples. The important limit in such a confidence interval is the upper limit. This upper limit is the P value that we may expect the P value of any given 13-trial sample to be below with a 97.5% probability. If this confidence interval upper limit is significant ($P < 0.05$), we may say that 13-trial samples for that network produce robustly significant changes in the descriptor during updating. The 95% CI upper limit P value for 13-trial samples using large saccades (45–50°) were as follows.

The max activation descriptor showed a robust, significant decrease only for the Vtop updating model: $P = 0.0000028$ (Vtop), $P = 0.26$ (Mtop), $P = 0.15$ (EV). The spread descriptor showed a robust, significant increase for the Mtop and EV models: $P = 0.15$ (Vtop), $P = 0.035$ (Mtop), $P = 0.019$ (EV). Note that the Vtop model failed to show a robust significant change in spread because the range of spread values across different trials increased to include both larger and smaller values (Fig. 8B, left, red). Finally, the lateral shift of activity showed a robust, significant increase for the Vtop and model alone: $P = 0.0016$ (Vtop), $P = 0.079$ (Mtop), $P = 0.063$ (EV).

For samples composed of 100 trials, *all* of the above-listed CI upper limit P values were significant ($P < 0.05$). This was because smaller differences in the means of two samples may be detected with larger sample sizes. However, in attempting to detect such fine differences, it is probable that the neural noise would become a problematic factor.

In summary, there were four robust and significant descriptor differences between the updating models that were robust across trial samples. First, was the speed and timing of the remapping as indicated in Fig. 8A, black, by the mean and SD fraction remap values. The Vtop model differed from the other two models in that its fraction remap was on average complete by the first updating time step (at 50 ms prior to saccade onset), whereas for the other models, it had only begun: *fraction*

[9] The last time step was chosen as the reference for the lateral shift descriptor because of the increase in this descriptor from the pre- to the post-updating time steps.

remap = 0.95 ± 0.35 (Vtop), 0.11 ± 0.12 (Mtop), 0.01 ± 0.02 (EV) For 13-trial samples, the *t*-test comparison of fraction remapping as of the first updating time step between the updating models, in terms of the above-described 95% CI upper limits were: $P = 0.00006$ (Vtop-Mtop), $P = 0.00001$ (Vtop-EV), $P = 0.055$ (Mtop-EV). Thus for 13-trial samples, the remapping as of the first updating time step (50 ms prior to saccade onset in our models) was robustly significantly different between the Vtop and other two models, somewhat less robustly significantly different between Mtop and EV models. When samples with larger numbers of trials were used, the difference in remapping were all robustly significant.

For the Mtop and EV models, remapping proceeded progressively from the start of the updating interval for the Mtop model and from saccade onset for the Vtop model. In both models, remapping completed only by the end of the saccade. As of the first saccade time step the values were: *fraction remap* = 1.06 ± 0.12 (Vtop), 0.65 ± 0.20 (Mtop), 0.23 ± 0.04 (EV). For 13-trial samples, the *t*-test comparison of fraction remapping between the updating models, in terms of the 95% CI upper limits, were: $P = 0.0037$ (Vtop-Mtop), $P = 10^{-13}$ (Vtop-EV), $P = 0.00001$ (Mtop-EV). Thus for 13-trial samples, the remapping as of saccade onset was robustly significantly different between all three updating models.

The second robust difference was that for the Vtop model, the maximum activation was strongly suppressed during the time steps immediately after the first updating time step (−50 ms), whereas in the Mtop and EV models there was no such significant suppression (the statistics of this being described in detail in the preceding text). The third robust difference was that there was a large increase in the lateral shift of the activity in the Vtop model prior to saccade onset and a smaller increase for the Mtop model that rose toward saccade onset and decreased during the saccade, whereas there was no significant increase for the EV model (statistics details given in the preceding text). The fourth robust difference was a robust and significant increase in spread of activity during remapping in the Mtop and EV models but not the Vtop model (details given in the preceding text).

In terms of using these differences in output-layer remapping to distinguish between updating models, we describe two practical approaches. The first is that already used by Sommer and Wurtz (2006): looking at the evolution of activity at the remapping midpoint. The second is that of detecting lateral-shift in activity.

SOMMER AND WURTZ TEST PARADIGM. Sommer and Wurtz (2006) recorded the evolution of activity of FEF neurons during updating across saccades. Updating vectors were chosen so that the receptive fields of these neurons were centered at the midpoint between initial and updated target positions—i.e., the remapping midpoint. We repeated this paradigm for our standard updating models and show the results in Fig. 9. The mean activity at the remapping midpoint in each time step for sets of 500 trials in each of five different saccade-size ranges (5–10, 15–20, 25–30, 35–40, 45–50°) is shown in Fig. 9*A*. The mean ± SD activities are shown for the small, medium, and large saccade-size sets in Fig. 9, *B–D*.

Any change in remapping midpoint activity during updating required that the activity at the midpoint changed from that before and after updating. Before updating, the output-layer target-position activity was a hill located at the initial target position, while after updating the hill was located at the updated target position. The remapping midpoint was defined as midway between these points so that the activity at this midpoint should be the same both prior to and after updating. This activity arose from the midpoint's location on the "shoulder" of the initial or final hill of activity. Because the hill of activity in the output layer was coherent prior to and after updating, this shoulder activity depended only on the distance of the midpoint from initial or updated target position. For perfect hills of activity located at the ideal initial or updated positions for 10, 20, 30, 40, and 50° saccades, shoulder activations at the midpoint were: 0.89, 0.64, 0.37, 0.17, and 0.06, respectively. Thus the larger the saccade the smaller the shoulder activity at the midpoint should be, as may be seen in Fig. 9, where the activity in the first five and last five time steps follows this pattern, and is similar for all three models.

The strong suppression of activity that was observed prior to saccade onset in the Vtop model (Fig. 8*A*, red) meant that the change in activity at the remapping midpoint during updating seen in Fig. 9 for the Vtop model showed a saccade size effect: a suppression that was more pronounced for smaller saccades. The lack of suppression observed in the Mtop and EV models during updating (Fig. 8*A*, red) meant that the change in activity during updating seen in Fig. 9 for these models also showed a saccade size effect: an increase in activity that was more pronounced for *large* saccades.

These changes in remapping midpoint activation during updating were robust and significant for 13-trial samples for all saccade sizes in the Vtop model but only for the largest trials for the Mtop and EV models.[10] The 95% CI upper limit *P* values for 13-trial samples were, for small saccades (5–10°): $P = 0.0003$ (Vtop), $P = 0.18$ (Mtop), $P = 0.58$ (EV); for medium saccades (25–30°): $P = 0.0007$ (Vtop), $P = 0.56$ (Mtop), $P = 0.43$ (EV); for large saccades (45–50°): $P = 0.004$ (Vtop), $P = 0.042$ (Mtop), $P = 0.000002$ (EV). And were robust and significant for 100-trial samples for all saccade sizes and models except for medium saccades in the Mtop model (95% CI upper limit $P = 0.149$).

DETECTING LATERAL SHIFT IN ACTIVITY. We define the *direct remapping corridor* as comprising all points that lie within a fixed distance from the direct remapping vector that connects initial and updated target positions. Activity outside this corridor refers to the activity of all output-layer units the preferred direction of which lies outside the corridor. The half-width of an ideal hill of activity in the output layer was 15°, and the curvature in remapping due to the torsional remapping occurring later than the nontorsional remapping was on average 4°. Thus for a coherent hill of activity with negligible lateral shift and spread increase, we expect only minor activity outside direct remapping corridors the half-widths of which are 20 or 30°. The increases in lateral shift and spread of activity during updating observed in the Vtop and Mtop models (Fig. 8*B*, black and red) produce transient activity that lies outside the direct remapping corridor from initial to updated target position in the output layer

---

[10] Note that the relevant saccade size dimension is its relation to the width of the neuron's receptive field. In our models, the units had 1/e widths of 15°, so the 10, 20, 30, 40, and 50° saccades corresponded to, respectively, 0.7, 1.3, 2.0, 2.3, and 2.7 times this width.
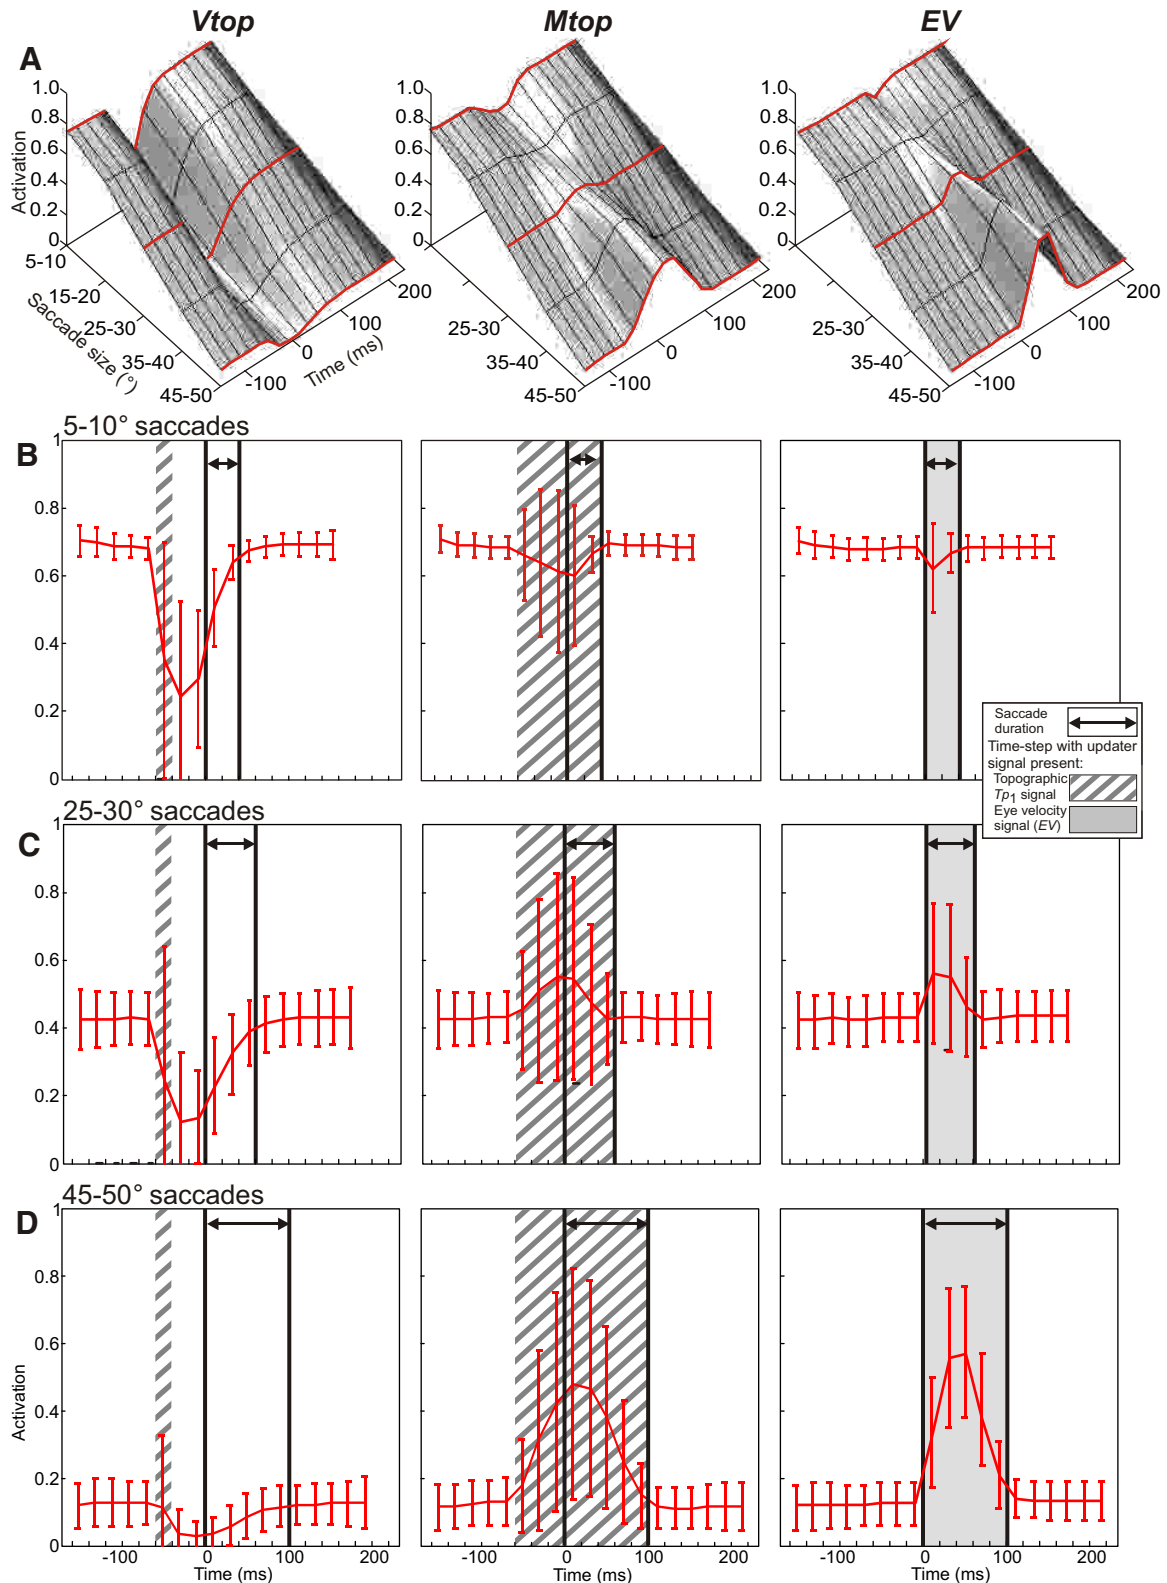
FIG. 9. Evolution of activity at the remapping midpoint during updating in the standard updating models. The activity at the remapping midpoint (the midpoint between the initial and updated 2nd-target positions) in each time step was calculated by interpolating the activations of the 4 adjacent output-layer topographic units, using a weighting that was the inverse of the distance to the units. *A*: the averages of this midpoint activity in each time step are plotted across sets of 500 trials of 5 saccade-size categories (5–10, 15–20, 25–30, 35–40, and 45–50°) for the Vtop, Mtop, and EV standard updating models. Also plotted are the means ± SD of this activity for small (5–10°; *B*), medium (25–30°; *C*), and large (45–50°; *D*) saccade trial sets.

of these networks. We illustrate this by plotting the maximum activity across all time steps for each output-layer unit in Fig. 10, *A–C*, for the three standard updating models for an example large-saccade updating trial. (The maximum activity is proportional to the size of the circle plotted at each output-layer unit's preferred direction.) The movement of the center of mass of activity in each model is shown by the dashed line. Remap corridors showing regions within 20 and 30° of the direct remapping vector (black line) are shown by the dark and light gray regions.

As may be seen, for this trial there is little activity outside the 20° corridor and none outside the 30° corridor for the EV model (Fig. 10*C*), while there is for both Vtop and Mtop models (*A* and *B*). The maximum activities in each time step outside the 20 and 30° remapping corridors are shown in blue and red for all models in Fig. 10, *D–F* (mean ± SD across a set of 500 trials). The sharp increase prior to saccade onset seen in the outside-corridor maximum activity for the Vtop model (Fig. 10*D*) and the smaller increase seen for the Mtop model (Fig. 10*E*) show that detection of activity at locations laterally displaced from the direct remapping vector (especially beyond the 30° lateral displacement, red in Fig. 10, *D–F*) may be used to distinguish the three updating models. For the EV model, there was no significant activity outside this 30° corridor during remapping

(Fig. 10*F*, red). For the Mtop model, there was significant activity outside this corridor during remapping, which occurred both prior to and during the saccade (Fig. 10*E*, red). For the Vtop model, the significant activity outside this corridor during remapping occurred only prior to saccade onset (Fig. 10*D*, red). Again using the 95% CI upper limit *P* values for 13-trial samples and using a remap corridor of radius 30° yielded significant increases in max activation outside the remap corridor for the Vtop and Mtop models only: $P = 0.0038$ (Vtop), $P = 0.032$ (Mtop), $P = 0.15$ (EV), For 100-trial samples, the 95% CI upper limit *P* values were significant for all models.

Results obtained from the variations of the updating models

The results described so far refer only to our three standard updating models. We also examined the effect of using variations of these networks. Brief summaries of the results of each variation are given here. The full description of these variation model results are given in the supplementary materials. Only the results that proved robust across all model variations are of interest. These results are treated with in the DISCUSSION.

EFFECT OF COMBINING TOPOGRAPHIC AND EV UPDATER SIGNALS. We trained two updating models in which the EV and a topographic updater signal (either Vtop or Mtop) were both used.
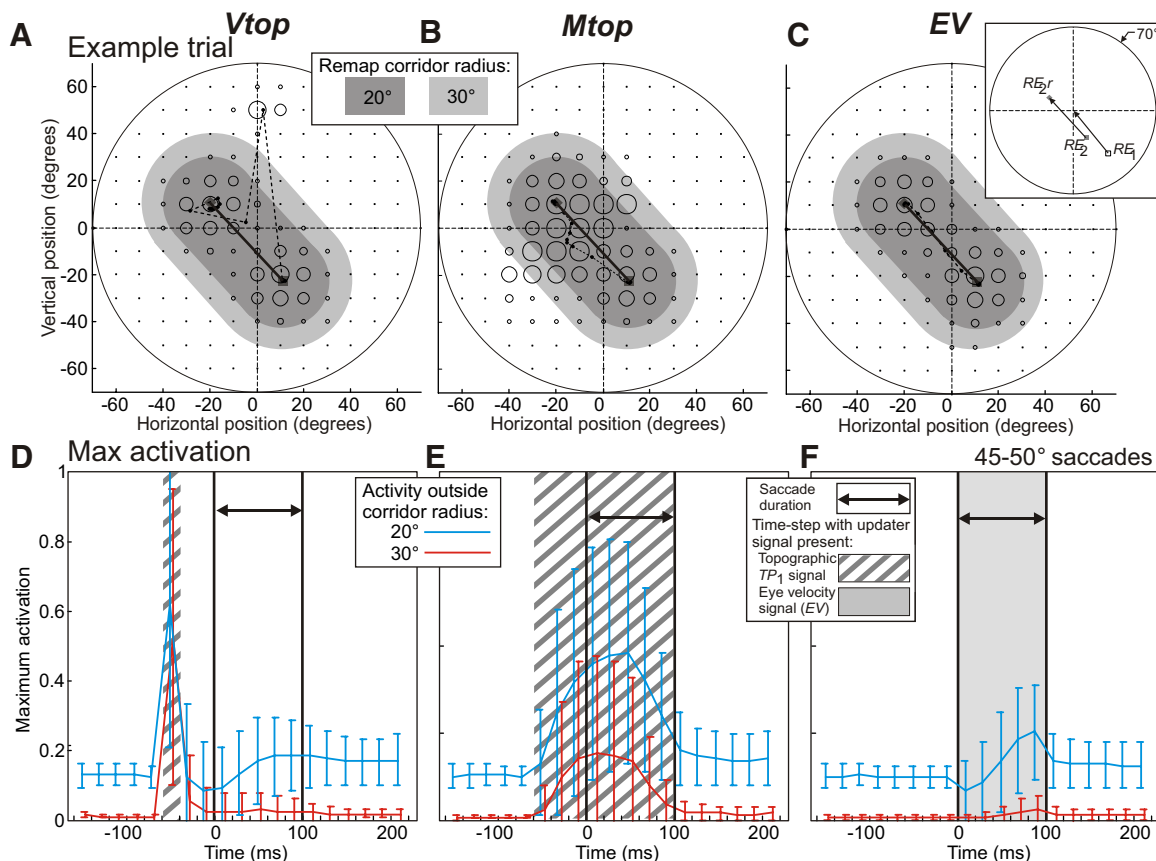

FIG. 10. Evolution of maximum activity outside the remapping corridor in the standard updating models. *A–C*: the maximum activations of each output-layer topographic unit, across all time steps for an example large-saccade trial (shown in *inset*) are plotted for the Vtop, Mtop, and EV standard updating models. This maximum activation is plotted as a black circle with diameter proportional to this activation, centered at the unit's preferred direction. The center of mass of activity in each time step is plotted as a small black dot, connected across time steps by a dashed line. The 20 and 30° remapping corridors, comprising the area within 20 and 30°, respectively, of the remapping vector (heavy black line) connecting the initial, TP₂, and updated, TP₂u, second-target positions, are shown as the dark and light gray regions. *D–F*: the maximum activity outside the 20 and 30° remapping corridors are plotted (blue and red, respectively) for each time step, mean ± SD over a set of 500 large-saccade (45–50°) trials for the Vtop, Mtop, and EV standard updating models.

Both of these *combined* updating models had remapping properties similar to the standard updating model using the associated topographic updater signal. The single exception to this occurred in the EV+Mtop updating model, where the suppression of activity during the saccade exceeded both EV and Mtop models. This exception was not important, however, because it will be shown that strong suppression in the unwarped topographic updating models does not persist in warped-field models. The evolutions of the four activation parameters for the combined models are shown for the largest saccade sizes in Supplementary Fig. S4. We may summarize our findings by saying that when both EV and a topographic updater signal are present, the latter signal dominates in determining remapping in the output layer.

EFFECT ON REMAPPING OF CHANGING TIME STEP DURATION. The choice of 20-ms time-step durations meant that saccades lasted from 1 to 5 time steps. As we have already stated, because the largest remapping differentiation between models occurred for the largest saccade sizes, we have compared updating trials that comprise remapping across 9 of these 20-ms time steps. To examine whether this time step duration produced artifacts in the remapping of our updating models we trained the three standard models using time steps that were decreased by a factor of 3, corresponding to 6.7 ms per time step.

In these short-time-step updating models, the large-saccade trials involved remapping across 27 of these 6.7-ms time steps. The output layer activation evolution, presented in Supplementary Fig. S8, showed descriptor behaviors similar to the standard updating models. The only significant difference was in the Vtop model where the recovery from suppression of activity that occurred in the time steps immediately following that in which the Vtop signal (Fig. 8*A,* red) was more rapid for the short-time-step model (Supplementary Fig. S8*A,* red). Because the transient suppression itself was preserved and because the timing onset of the Vtop signal was arbitrary, this recovery timing difference was not included as one of the remapping features used to differentiate between updating models.

EFFECT OF THE SPATIAL WARPING ON THE REMAPPING. In our standard updating models, we did not represent any warping in the topographic representation of second-target position in the input or output layers of our networks. To test the effect such warping might have on remapping, we trained a neural network to perform updating within a hemifield topographic map that was warped in a manner similar to that observed in the SC (supplementary materials). These networks showed remapping patterns that were substantially similar to the standard updating models, with some minor variations that are detailed in the supplementary materials. These differences will be addressed in the DISCUSSION.

DISCUSSION

In contrast to our previous theoretical studies (Keith and Crawford 2008; Keith et al. 2007), which focused on the geometric aspects of potential remapping mechanisms alone, the current study explored both the spatial and temporal properties of remapping. We trained recurrent neural networks to perform target-position updating across saccades and examined how the spatiotemporal properties of the signals used to drive the updating (updater signals) affected the associated remap-

ping in both the network's output layer and the upstream hidden layer. We placed no constraints on how remapping occurred. Our analysis focused the *emergent properties* of these networks: remapping in the topographic output layer proceeded and the mechanism underlying this in the hidden layer. Our goal was to provide a theoretical framework that would allow physiologists to ask questions such as: is there just one mechanism for spatial updating or are their multiple mechanisms? Are different mechanisms associated with different brain areas? Is each neuron within one area driven by the same updater signals? Can one distinguish between these signals by recording their effects? Toward answering these questions, we will review the major features and predictions of our models and then relate these to the known physiology.

Remapping in the output layer: findings and predictions

The output layer of the network was a major focus of this study because this layer represented the 2-D topographic brain areas that show the *results* of remapping—i.e., the signals that physiologists are most likely to record during experiments. Because this is the first study to address dynamics and topography of remapping in a recurrent network, we had little to guide us besides intuition and known physiology. For the sake of generalization, we simulated a generic topography in the main results and figures, but for the sake of comparison, we also did these simulations with the warped topography observed in the SC (as reported in the supplementary materials). We will consider both versions here. The effect of warping did not change the basic differences in remapping patterns that distinguished our updating models: the progress of remapping (including the degree of presaccadic remapping), also the spread and displacement of activity during remapping, in the topographic output layer; and the degree and temporal evolution of the spread of remapping across the population of units comprising the hidden layer.

The consistent and robust differences among the three updater models we considered (Vtop, Mtop, and EV) provide several specific testable predictions between these models. Most basically, when updating was driven by eye velocity alone (the EV updating model), the remapping of this activity took the form of a coherent moving hill that only began at the start of the saccade (Fig. 8 and Supplementary Fig. S10). In contrast, when remapping was driven by topographic predictive measures of the saccade target (Vtop and Mtop, representing the visual or motor activity in FEF or SC), presaccadic remapping was produced—on average full remapping for Vtop and partial remapping for Mtop (Fig. 8*A* and Supplementary Fig. S10*A*). Moreover, for the Mtop model, the hill of activity was less coherent, more spread during the remapping (Fig. 8*B* and Supplementary Fig. S10, *B–D*); and for the Vtop model, the activity was presaccadically highly sporadic both in distribution and position (Fig. 8*B* and Supplementary Fig. S10, *B–D*) and jumped rather than moved to the updated target position (Fig. 8*A* and Supplementary Fig. S10*A*), gradually settling into a coherent hill of activity at the updated target position by the end of the saccade (Fig. 8, *A* and *B,* and Supplementary Fig. S10).[11]

---

[11] While the timing of this sporadic dispersion of activity in our model did not encompass the saccade itself, the effect itself is consistent with the transient lack of remapped activity observed in the SC by Walker et al. 1995.

The gradually moving activity associated with the EV and Mtop models (coherent hill for the EV model and more spread distribution for the Mtop model) predict that a significant amount of activity will pass through the remapping midpoint studied by Sommer and Wurtz (2006) in FEF, while in the Vtop model sporadic, jumping distribution does not (Fig. 10 and Supplementary Fig. S12). Also the *change* in activity at the midpoint during remapping relative to the pre- and postremapping period activity shows a distinct saccade-size effect. This is also complicated by topographic field warping because the receptive field sizes tend to follow this warping (e.g., Cynader and Berman 1972), thus making the shoulder activity at the midpoint seen in the pre- and postremapping periods dependent on this warped-field remapping distance rather than saccade size alone. Thus when the remapping is such that the shoulder activity is large (small warped-field remapping distance) the Vtop model predicts a strong suppression in midpoint activity during remapping, while the EV and Mtop models predict no significant suppression. For remapping with small shoulder activity (large warped-field remapping distance), on the other hand, the EV and Mtop models predict a large increase in the midpoint activity during remapping, while the Vtop model predicts no such increase (Fig. 10C and Supplementary Fig. S12).

Although both EV and Mtop models predict some lateral displacement of activity during remapping, only the Mtop model predicts a significant increase in spread of activity (Fig. 8 and Supplementary Fig. S10). The detection of the increase (or lack thereof) of activity spread would require a use of multielectrode recording in the topographically organized structure being studied. By varying the updating vector, the same set of electrodes may record simultaneous activity corresponding to different sets of positions relative to the direct updating vector.[12] In this way, both spread and displacement of activity may be measured (although the latter does not require the use of multiple electrodes). The presence of the presaccadic sporadic activity predicted by the Vtop model may be detected by looking for transient activity prior to saccade onset at positions far from the direct remapping vector.

The preceding represent the robust findings of our updating network simulations, and their associated predictions, in terms of the topographic structure in which target position updating is represented by activity remapping across saccades. These findings are robust in that they were consistent across the three principal variations of our updating networks: the standard model networks, those with a threefold shorter time step, and those incorporating the warping observed in the SC. We expect that the common method of averaging values across multiple trials used in neurophysiological studies (e.g., Sommer and Wurtz 2006) will reduce the standard errors of the mean to significantly different values for all robust results reported here.

These various predictions suggest that a major goal of our study is feasible: one should be able to distinguish the types of signals that drives updating from the dynamics of remapping. Note, however, that these various predictions might distinguish the model that best explains a given set of unit properties in a given area of the brain, but any such dataset should not be used

to infer that all other neurons and brain areas use the same mechanism. It is quite possible, and even likely, that the real brain combines these signals in an optimal fashion.

Mechanisms underlying remapping: hidden layer properties

Our analysis of the hidden layer suggested that in general, target position was sustained by means of the recurrent connections that sent excitatory feedbacks from each HLU to those HLUs having similarly positioned receptive fields, and inhibitory feedbacks to those HLUs having highly differently positioned receptive fields. This mechanism was first described in recurrent neural networks by Xing and Andersen (2000).

As we have seen, remapping of target positions in the output layer was highly influenced by the updater signals, which in turn could be represented as operating through gain-field modulations of hidden-layer activations. The total updater-signal input to a given HLU, plotted as a function of saccade-target position, yielded what we termed the updater input field to that unit. The updater input fields for the Vtop and Mtop models were highly irregular, whereas those for the EV model were highly regular (Fig. 4). The most important effect of this distinction was that for the EV model the transient remapping of HLU receptive fields was quite regular as well, such that the spread across all HLUs was small in any time step (Fig. 3C). The remapping in both the Vtop and Mtop models showed a large transient spreads of remapping across their HLUs (Fig. 3, A and B). This spread was limited to presaccadic time steps for the Vtop model, whereas for the Mtop model, it increased toward saccade onset and then decreased during the saccade. The different patterns of remapping spread in the three updating models were robust even when spatial warping was introduced (Supplementary Fig. S13). Thus there was a direct correlation between coherent HLU remapping and the maintenance of a coherent hill in the output layer as seen in the *EV* model. The temporal evolution of the spread of remapping across HLUs followed that of the degree of noncoherence in the activity in the output layer. This reflected the fact that it was in the hidden layer that the calculations generating output-layer activity were performed. The output layer was merely a read-out of the state of the hidden-layer activity.

To sum up, we were able to both identify the mechanisms for updating in our networks and show that these mechanisms and their effects lead to different experimental predictions for both the hidden layer and the output layer [for example as shown in Fig. 10 and Supplementary Fig. S12, using the experimental paradigm devised by Sommer and Wurtz (2006)]. Many of these predictions have not yet been tested, but some can be compared directly to the known physiology.

On the dependence of remapping on time-step duration

The patterns of remapping observed in networks trained using the same updating model (i.e., Vtop, Mtop, or EV) but with 20- and 6.7-ms time-step durations (standard and short-time-step networks), were remarkably similar (Fig. 8 and Supplementary Fig. S8). The single significant difference in these patterns occurred for the Vtop updating model. In this model, the transient Vtop updater signal lasted only a single time step irrespective of the time-step duration. The recovery of both the maximum activation and spread of activation (Fig.

---

[12] By positions, we mean, of course, the locations of the neuron receptive fields. The resolution of position is thus limited by the width of these receptive fields.

8*A*/Supplementary Fig. S8*A, left*, red, and Fig. 8*B*/Supplementary Fig. S8*B, left*, red) was better described in terms of the number of time steps than the time interval. The principal effect of this was that the activation recovered prior to saccade onset in the short-time-step Vtop network but not in the standard Vtop network. However, we noted in METHODS that the onset of the Vtop signal chosen in these networks was arbitrary, so that the degree of activation recovery in the Vtop model as of saccade onset could not be considered a reliable result of these simulations.

In all other respects, a comparison of the remap descriptors for the standard and short-time-step networks within each model showed a strikingly good similarity (bearing in mind that some difference is expected between any two networks because each represents a separate solution of the updating problem). This similarity shows that the remapping is determined by the evolution of updater signals, which (other than the Vtop signal) varied fairly smoothly across the duration of the saccade and updating interval. This may be derived from the fact that the time scale of the recurrent connections in the hidden layer by which the activity was sustained, assuming a single interneuron in the feedback loop are of the order of the short-time-step duration (~ 5 ms). Yet the use of such a time step versus using one three times as long (the standard, 20-ms time step, duration), produced, again with the single exception of the recovery time of activation and spread in the Vtop model, the same remapping patterns. This in turn produces (again with the same exception) the similar patterns of activity evolution at the remapping midpoint for networks trained with the same model but different time steps (Fig. 0/Supplementary Fig. S9). This illustrates the point that the remapping properties in updating networks depend not on the time-scale of connections such as the recurrent connections within the hidden layer but on the time scale of the evolution of signals used to drive the updating.

Possible physiological analogues of the hidden layer

The properties of the hidden layer of our updating networks emerged from the training of these networks. There was no a priori topographic structure assigned these units. The properties they developed were: activity latching while a target position was being held, activity remapping during updating across saccades, and a transient spread of remapping across units during this updating for the Vtop and Mtop updating models. The first two properties are those known to be associated with area LIP neurons (Barash et al. 1991; Colby et al. 1996; Duhamel et al. 1992). The third property has also been observed (in preliminary reports) in area LIP but not VIP neurons during updating across saccades (Kubischik and Bremmer 1999). These observations suggest that the hidden layer of our networks might best be identified with area LIP. A final consideration regards the presence of open and closed receptive fields in our HLUs.

All three updating model produced both open (including sigmoidal) and closed receptive fields in their HLUs. While our recurrent updating networks predominantly developed HLUs with open receptive fields, this was not a requirement mandated by the limitation of open-field-only remapping that exists in feed-forward networks but merely arose from the efficiency of such fields in performing the required task of remapping in

the output layer. As in any neural network, our models represented a greatly simplified picture of the enormous populations of neurons and the complex properties of their various connections. For example, in this network we did not seek to include the rich set of neuron types observed in structures such as the SC and FEF (Bruce and Goldberg 1985; Schall 1991). This was because the goal of our networks was limited to simulating fundamental computational aspects of remapping rather than intentionally reproducing the full physiology and anatomy associated with saccade initiation and control (Heinzle et al. 2007). The advantage of this approach is that any resulting similarities with physiology can be attributed to fundamental computational constraints rather than contrived design.

If closed receptive fields are required to dominate in the hidden layer, additional constraints would have to be imposed, such as those used by Mitchel and Zipser (2001) in their updating models. But such techniques themselves introduce constraints that are unconfirmed in terms of their relation to real biological brain systems, such as requiring all connections from the hidden to output layer to be excitatory only. We note that apart from differences in their relative numbers, our updating models and neurons in area LIP are similar in that they display both open and closed receptive fields in addition to other properties. The important point is that our models are *able* to perform full closed-field remapping, such as is seen in area LIP neurons.

These properties appear to identify LIP with the hidden unit layer of our model and conversely implicate LIP in the mechanism of remapping. However, the stronger claim that LIP is *the* mechanism for remapping (and other areas just report the result) would be premature. The real brain has anatomic and biological constraints not present in our model, including the apparent need to distribute saccade control and remapping over a number of different areas. And each of these areas communicates directly or indirectly with the other (Andersen et al. 1990; Ferraina et al. 2002; Lynch et al. 1985; Schall et al. 1995; Stanton et al. 1995). Thus the possibility remains that each of these areas (occipital cortex, LIP, SEF, FEF, and SC) is involved in the mechanism or mechanisms of remapping. Only more data can answer this question.

Possible physiological analogues of the output layer

Because this study raises the possibility that different brain areas and even different neurons within a given area might be driven by different updater signals, this question can only be answered by a careful comparison across neurons and brain areas. To date the most complete study of intrasaccadic remapping dynamics was done by Sommer and Wurtz (2006) on the FEF. They used the method (simulated here in Fig. 10 and Supplementary Fig. S12) of recording the activity neurons the receptive fields of which were centered at the midpoint of remapping to distinguish the moving or spreading from jumping hill models. We were able to replicate the main aspects of their findings (no significant increase or decrease of activity during remapping across medium-sized saccades)[13] with the Mtop version of our model (which also reproduced presaccadic

---

[13] The "shoulder" activity reported in Sommer and Wurtz (2006) was roughly half the presaccadic activity at the initial target position. For our standard updating models, this corresponded to saccade sizes of $\sim 30°$, which were our medium-sized saccades.

remapping). This finding squares well with the experimental work of Sommer and Wurtz (2004a,b, 2006, 2008) suggesting that the superior colliculus motor output provides the updater signal to the cortex via the thalamus. However, our simulations also point out the difficulty of interpreting such data. For example, whereas they tested between simple jumping hill and moving hill models, our simulations suggest that other, perhaps unexpected possibilities, such as suppressed moving hills and dispersed activity are possible. Moreover, we found there to be a strong effect of saccade size on the pattern of midpoint activity observed during updating (due to the fact that the shoulder activity registered at the midpoint when a hill of activity is centered at either the initial or updated target position before and after updating decreases as saccade size increases). Our simulations suggest that further information can be gleaned by doing such experiments with different saccade sizes and testing points off the main line of updating.

Other areas of the brain have not been tested in the same way, but there are several clues. The marked suppression/dispersion of remapped activity in the SC during saccades appears to be consistent *in principle* with our Vtop model if the time constant of the suppression is extended to include the entire saccade, but again this would need to be checked with different saccade sizes to see if the suppression already observed (Walker et al. 1995) is generally observed. It is also possible that the SC has intrinsic inhibitory connections for other purposes (Munoz and Istvan 1998) that would provide the same effect, so other predictions of the model would need to be directly tested. Presaccadic remapped activity has been observed in SC, FEF, and LIP (Duhamel et al. 1992; Umeno and Goldberg 1997; Walker et al. 1995). This can be explained by our Vtop and Mtop models, but is not consistent with the EV variation of our model. However, not all neurons even in these areas show presaccadic remapping. Thus it is tempting to hypothesize that these areas are driven by a variety of signals, including those resembling Vtop, Mtop, and EV. The only way to test this would be to repeat some of the experiments simulated here and cross-correlate whether units that show/do not show presaccadic remapping also show other properties consistent with the different variations of our models.

The best available report of a direct comparison of remapping in different visuomotor areas looked at passive remapping of visual responses during saccades in areas MIP, VIP, and MT/MST (Bremmer et al. 2009; Kubischik and Bremmer 1999). Consistent with our basic premise, these preliminary reports suggest different patterns of remapping observed in different brain areas. For example, Kubischik and Bremmer (1999) described neurons in area LIP that showed a spread of remapping matching our Mtop updating model: a wide, transient spread of remapping prior to and during the saccades. VIP and MT/MST did not appear to show the same results.

Conclusions

The general goal of this study was to provide a theoretical framework for future unit recording studies of saccade-related remapping. More specifically, our aim of this study was to determine if the type of saccade efference copy signal used for spatial updating influences the mechanism and dynamics of remapping and if these differences should be detectable through the use of existing or plausible experimental paradigms. In this we believe we were successful. We found that a neural network model trained on three different types of updating signal (Vtop, Mtop, and EV) produced three quite different patterns of peri- and intrasaccadic remapping, both in the hidden layer (the mechanism) and the output layer (the result). As described in detail in the preceding text, some of the properties we observed have already been observed in several areas of the brain, but most have not yet been tested. Especially important would be a comparison across brain sites and neuron classes involved in remapping. Moreover, this study provides a cautionary tale for physiologists: even a simple network can show emergent properties far more complex than the simple moving hill versus jumping hill dichotomy that has guided previous experiments.

REFERENCES

Andersen RA, Asanuma C, Essick G, Siegel RM. Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J Comp Neurol* 296: 65–113, 1990.

Andersen RA, Essick GK, Siegel RM. Encoding of spatial location by posterior parietal neurons. *Science* 230: 456–458, 1985.

Baker JT, Harper TM, Snyder LH. Spatial memory following shifts of gaze. I. Saccades to memorized world-fixed and gaze-fixed targets. *J Neurophysiol* 89: 2564–2576, 2003.

Barash S, Bracewell RM, Fogassi L, Gnadt JW, Andersen RA. Saccade related activity in the lateral intraparietal area. I. Temporal properties, comparison with area 7a. *J Neurophysiol* 66: 1095–1108, 1991.

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

Bremmer F, Kubischik M, Hoffman K-P, Krekelberg B. Neural dynamics of saccadic suppression. *J Neurosci* 29: 12374–12383, 2009.

Bruce DJ, Goldberg ME. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* 53: 603–635, 1985.

Cai RH, Pouget A, Schlag-Rey M, Schlag J. Perceived geometrical relationships affected by eye-movement signals. *Nature* 386: 601–604, 1997.

Cassanello CR, Ferrera VP. Computing vector differences using a gain field-like mechanism in monkey frontal eye field. *J Physiol* 582.2: 647–664, 2007.

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

Colby CL, Duhamel JR, Goldberg ME. Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J Neurophysiol* 76: 2841–2852, 1996.

Crawford JD. The oculomotor neural integrator uses a behavior-related coordinate system. *J Neurosci* 14: 6911–6923, 1994.

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

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

Cynader M, Berman N. Receptive-field organization of monkey superior colliculus. *J Neurophysiol* 35: 187–201, 1972.

Droulez J, Berthoz A. A neural network model of sensoritopic maps with predictive short-term memory properties. *Proc Natl Acad Sci USA* 88: 9653–9657, 1991.

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

Ferraina S, Pare M, Wurtz RH. Comparison of cortico-cortical and cortico-collicular signals for the generation of saccadic eye movements. *J Neurophysiol* 87: 845–858, 2002.

Hallett PE, Lightstone AD. Saccadic eye movements to flashed targets. *Vision Res* 16: 107–114, 1976.

Heinzle J, Hepp K, Martin KAC. A microcircuit model of the frontal eye fields. *J Neurosci* 27: 9341–9353, 2007.

Henn V, Hepp K, Vilis T. Rapid eye movement generation in the primate: physiology pathophysiology, and clinical implications. *Rev Neurol* 145: 540–545, 1989.

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

Herter TM, Guitton D. Human head-free gaze saccades to targets flashed before gaze-pursuit are spatially accurate. *J Neurophysiol* 80: 2785–2789, 1998.

Honda H. Perceptual localization of visual stimuli flashed during saccades. *Percept. Psychophys.* 45: 162–174, 1989.

Keith GP, Crawford JD. Saccade-related remapping of target representations between topographic maps: a neural network study. *J Comput Neurosci* 24: 157–178, 2008.

Keith GP, Smith MA, Crawford JD. Functional organization within a neural network trained to update target representations across 3-D saccades. *J Comput Neurosci* 22: 191–209, 2007.

Keller EL, Edelman JA. Use of interrupted saccade paradigm to study spatial and temporal dynamics of saccadic burst cells in superior colliculus in monkey. *J Neurophysiol* 72: 2754–2770, 1994.

Klier EM, Wang H, Crawford JD. Three-dimensional eye-head coordination is implemented downstream from the superior colliculus. *J Neurophysiol* 89: 2839–2853, 2003.

Kubischik M, Bremmer F. Perisaccadic space representation in monkey inferior parietal cortex. *Soc Neurosci Abstr* 25: 1164, 1999.

Luschei ES, Fuchs AF. Activity of brain stem neurons during eye movements of alert monkeys. *J Neurophysiol* 35: 445–461, 1972.

Lynch JC, Graybiel AM, Lobeck LJ. The differential projection of two cyto-architectonic subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus. *J Comp Neurol* 235: 241–254, 1985.

Lynch JC, Tian JR. Cortico-cortical networks and cortico-subcortical loops for the higher control of eye movements. *Prog Brain Res* 151: 461–501, 2005.

Matin L, Matin E, Pearce DG. Visual perception of direction when voluntary saccades occur. I. Relation of visual direction of a fixation target extinguished before a saccade to a flash presented during the saccade. *Percept Psychophys* 5: 65–78, 1969.

Mays LE, Sparks DL. Dissociation of visual and saccade-related responses in superior colliculus neurons. *J Neurophysiol* 43: 207–232, 1980.

Medendorp WP, Smith MA, Tweed DB, Crawford JD. Rotational remapping in human spatial memory during eye and head motion. *J Neurosci* 196: 1–4, 2002.

Medendorp WP, Tweed DB, Crawford JD. Motion parallax is computed in the updating of human spatial memory. *J Neurosci* 23: 8135–8142, 2003.

Miller JM. Information used by the perceptual and oculomotor systems regarding the amplitude of saccadic and pursuit eye movements. *Vision Res* 20: 59–68, 1980.

Mitchell J, Zipser D. A model of visual-spatial memory across saccades. *Vision Res* 41: 1575–1592, 2001.

Munoz DP, Istvan PJ. Lateral inhibitory interactions in the intermediate layers of the monkey superior colliculus. *J Neurophysiol* 79: 1193–1209, 1998.

Munoz DP, Pelisson D, Guitton D. Movement of neural activity on the superior colliculus motor map during gaze shifts. *Science* 251: 1358–1360, 1991.

Munoz DP, Wurtz RH. Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and build-up neurons. *J Neurophysiol* 73: 2313–2333, 1995.

Ohtsuka K, Noda H. Burst discharges of mossy fibers in the oculomotor vermis of macaque monkeys during saccadic eye movements. *Neurosci Res* 15: 102–114, 1992.

Port NL, Sommer MA, Wurtz RH. Multielectrode evidence for spreading activity across the superior colliculus movement map. *J Neurophysiol* 84: 344–357, 2000.

Quaia C, Optican LM, Goldberg ME. The maintenance of spatial accuracy by the perisaccadic remapping of visual receptive fields. *Neural Networks* 11: 1229–1240, 1998.

Riedmiller M, Braun H. A direct adaptive method for faster backpropagation learning: the RPROP algorithm. *IEEE Int Conf on Neur Netw*: 586, 1993.

Ross J, Morrone MC, Burr DC. Compression of visual space before saccades. *Nature* 386: 598–601, 1997.

Russo GS, Bruce CJ. Neurons in the supplementary eye field of rhesus monkeys code visual targets and saccadic eye movements in an oculocentric coordinate system. *J Neurophysiol* 76: 825–848, 1996.

Schall JD. Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: comparison with supplementary eye fields. *J Neurophysiol* 66: 559–676, 1991.

Schall JD, Morel A, King DJ, Bullier J. Topography of visual cortex connections with frontal eye field in macaque: convergence and segregation of processing streams. *J Neurosci* 15: 4464–4487, 1995.

Schiller PH, Sandel JH. Interactions between visually and electrically elicited saccades before and after superior colliculus and frontal eye field ablations in the rhesus monkey. *Exp Brain Res* 49: 381–392, 1983.

Schlag J, Schlag-Rey M, Dassonville P. Saccades can be aimed at the spatial location of targets flashed during pursuit. *J Neurophysiol* 64: 575–581, 1990.

Schlag-Rey M, Schlag J, Shook B. Interactions between natural and electrically evoked saccades. I. Differences between sites carrying retinal error and motor error signals in monkey superior colliculus. *Exp Brain Res* 76: 537–547, 1989.

Smith MA, Crawford JD. Implications of ocular kinematics for the internal updating of visual space. *J Neurophysiol* 86: 2112–2127, 2001.

Smith MA, Crawford JD. A distributed population mechanism for the 3-D oculomotor reference frame transformation. *J Neurophysiol* 93: 1742–1761, 2005.

Soetedjo R, Kaneko CRS, Fuchs AF. Evidence against a moving hill in the superior colliculus during saccadic eye movements in the monkey. *J Neurophysiol* 87: 2778–2789, 2002.

Sommer MA, Wurtz RH. What the brain stem tells the frontal cortex. I. Oculomotor signals sent from superior colliculus to frontal eye field via mediodorsal thalamus. *J Neurophysiol* 91: 1381–1402, 2004a.

Sommer MA, Wurtz RH. What the brain stem tells the frontal cortex. II. Role of the SC-MD-FEF pathway in corollary discharge. *J Neurophysiol* 91: 1403–1423, 2004b.

Sommer MA, Wurtz RH. Influence of the thalamus on spatial visual processing in frontal cortex. *Nature* 444: 374–377, 2006.

Sommer MA, Wurtz RH. Brain circuits for the internal monitoring of movements. *Annu Rev Neurosci* 31: 317–338, 2008.

Sparks DL. Neural cartography: sensory and motor maps in the superior colliculus. *Brain Behav Evol* 31: 49–56, 1988.

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

Sparks DL, Porter JD. Spatial localization of saccade targets. II. Activity of superior colliculus neurons preceding compensatory saccades. *J Neurophysiol* 49: 64–75, 1983.

Stanton GB, Bruce CJ, Goldberg ME. Topography of projections to posterior cortical areas from the macaque frontal eye fields. *J Comp Neurol* 353: 291–305, 1995.

Suzuki Y, Buttner-Ennever JA, Straumann D, Hepp K, Hess BJM, Henn V. Deficits in torsional and vertical rapid eye movements and shift of Listing's plane after uni- and bilateral lesions of the rostral interstitial nucleus of the medial longitudinal fasciculus. *Exp Brain Res* 106: 215–232, 1995.

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

Umeno MM, Goldberg ME. Spatial processing in the monkey frontal eye field. I. Predictive visual responses. *J Neurophysiol* 78: 1373–1383, 1997.

van Gisbergen JA, Robinson DA, Gielen S. A quantitative analysis of generation of saccadic eye movements by burst neurons. *J Neurophysiol* 45: 417–442, 1981.

van Opstal AJ, Hepp K, Hess BJM, Straumann D, Volker H. Two- rather than three-dimensional representation of saccades in monkey superior colliculus. *Science* 252: 1313–1315, 1991.

Walker MF, FitzGibbon EJ, Goldberg ME. Neurons in the monkey superior colliculus predict the visual result of impending saccadic eye movements. *J Neurophysiol* 73: 1988–2003, 1995.

White RL, Snyder LH. A neural network model of flexible spatial updating. *J Neurophysiol* 91: 1608–1619, 2004.

Xing J, Andersen RA. Memory activity of LIP neurons for sequential eye movements simulated with neural networks. *J Neurophysiol* 84: 651–665, 2000.

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