Neuron, Vol. 21, 661–664, October, 1998, Copyright ©1998 by Cell Press

Reading Neural Representations

Minireview

Jennifer M. Groh*
Center for Cognitive Neuroscience
Department of Psychology
Dartmouth College
Hanover, New Hampshire 03755

Neurons carry signals that are correlated with external sensory events, internal mental states, impending behavioral responses, and many different combinations of these factors. Single unit recording studies in awake animals have painted an increasingly thorough portrait of the types of signals present in different brain areas. Such studies reveal what information various populations of neurons encode, but they fall short of describing how this information is used. Neurons are, after all, connected to other neurons. How does each successive population of neurons interpret the input it receives from the preceding stage of processing and produce an output that will eventually manifest itself as a behavioral response? A thorough understanding of how the brain works will ultimately require a detailed characterization of how each neural representation is "read out" by subsequent stages of processing.

Progress in this challenging task requires manipulating the information present in neural representations by either removing or adding signals. Removal of neural signals can be accomplished by temporary pharmacological inactivation of neural tissue or by permanent surgical or chemical lesions. Such experiments have a rich history and can provide irrefutable evidence of the overall function served by the brain area under study, but lesions do not always reveal the specific details of that function. In contrast, the technique of microstimulation is capable of actually introducing neural signals at a specific location in the brain in a precisely controlled fashion. These artificially induced signals can help reveal the functional details of how the information in these areas is extracted by other areas.

In this article, I will describe how microstimulation has been employed to explore the algorithms used for reading neural representations. These studies come from two different areas of the brain, extrastriate cortical area MT, which contains a representation for visual motion, and the superior colliculus (SC), which contains a representation for the direction and amplitude of saccadic eye movements. Both areas use a "place code" to represent their respective types of information. MT neurons have receptive fields tuned for the retinotopic location of visual stimuli, and neurons are organized topographically according to the locations of their receptive fields. This forms a map for the position of stimuli in the visual scene. At a finer level of detail, MT neurons are sensitive not only to the location of visual stimuli but also to the direction and speed of moving stimuli. Their tuning for direction and speed can be thought of as producing a receptive field in the velocity domain in addition to the receptive field in the position domain. Local topography for the velocity receptive fields is superimposed on the global topography for the positional receptive fields: neurons sharing the same positional receptive field and whose velocity receptive fields have the same preferred direction of visual motion are clustered together to form cortical columns analogous to the ocular dominance and orientation columns in primary visual cortex (see references in Groh et al., 1997).

Like neurons in MT and many other sensory areas, SC neurons are tuned for the location of sensory stimuli. Because their activity is also correlated with impending saccadic eye movements, these neurons are sometimes said to have movement fields rather than receptive fields. The discharge of a typical SC neuron signals the direction and amplitude of a saccade to the location of a sensory stimulus at the center of its movement field. Again, the neurons are organized topographically so that their movement fields form a map for the vector of impending saccades, a map that is similar in flavor to the retinotopic maps so familiar in early stages of visual processing (see references in Stanford et al., 1996).

How are these representations read out by subsequent representations? The answer to this question requires some further discussion of the nature of the neural codes downstream. For the SC, this answer is relatively clear. Because the SC lies fairly close to the motor periphery, its efferent connections are reasonably well characterized, and the representational formats employed by the areas downstream from the SC are fairly well understood. The SC's main function is to provide command signals for eliciting saccadic eye movements. These commands are sent to the horizontal and vertical gaze control centers in the brainstem and ultimately to the motor neurons for the extraocular muscles (reviewed by Sparks and Hartwich-Young, 1989). The motor neurons, however, do not use a place code for signaling the desired eye movement. They do not have movement fields in the same sense that SC neurons do. Rather than being tuned for a particular movement vector, their activity is tuned for the pulling direction of their particular extraocular muscle, and the discharge rate is monotonically related to the position, speed, and acceleration of the eyes along that direction (reviewed by Carpenter, 1988). The read-out of the collicular motor map, then, must entail conversion of the place-coded signal for saccade vector into a firing rate code in which a low discharge rate produces a smaller amplitude movement and a higher discharge rate produces a larger amplitude movement.

For area MT, the situation is more complex. MT is a sensory area, and the motion information in this representation is likely to be used to guide many different types of behavioral responses, including smooth pursuit eye movements, movements of the body to intercept (e.g., catch a frisbee) or avoid a moving object (e.g., duck a punch), as well as perceptual judgements of motion direction and speed. Thus, the read-out mechanism is likely to depend on the type of behavioral response involved. However, for at least one of these behavioral responses, namely smooth pursuit eye movements, the constraints on the read-out mechanism are

*E-mail: groh@tahoe.cs.dartmouth.edu.

strikingly similar to those for the SC. Smooth pursuit eye movements, like saccades, are ultimately controlled by motor neurons that are tuned for the direction of the eye movement and have a firing rate that signals position, speed, and acceleration monotonically. Thus, MT's role in guiding smooth pursuit eye movements shares important characteristics with the SC's role in guiding saccades.

How might signals from a map-like representation such as the ones in the SC or MT be converted into the rate code employed by these motor systems? This requires an algorithm for converting the locus of activity in the map into a firing rate that corresponds to the preferred parameter value (i.e., saccade vector in the SC or visual stimulus velocity in MT) of the active neurons. For both the SC and MT, the signal encoded by the representation is a vector quantity, not a scalar, so the following discussion actually applies to the horizontal or vertical components of these vectors.

At least three mechanisms could potentially accomplish this task: weighted summation, weighted averaging, and winner-take-all. Under summation and averaging, all the neurons in the place code can "vote" for their preferred parameter value. The strength of each neuron's vote corresponds to the product of its activity level and its synaptic weight. Neurons that prefer larger parameter values (i.e., faster visual motion, larger amplitude saccades) have stronger synaptic weights than neurons that prefer smaller parameter values. The weighted votes are then either summed or averaged to produce an output value that is monotonically related to the parameter value. Winner-take-all involves a qualitatively different algorithm: neurons compete with one another, and the neuron (or small number of neurons) with the highest firing rate "wins." The activity of all other neurons in the place code is ignored, and the preferred parameter value of the winner becomes the output value.

Any of these methods can produce an output indicating the actual parameter value in the form of a rate code and is therefore potentially suitable for use as a motor command. However, these algorithms work best when only one parameter value is encoded in that map. The algorithms can be distinguished from one another by how they respond when more than one parameter value is present. When two parameter values are encoded in the place code, a summation mechanism will produce an output corresponding to the sum of the outputs that occur when either parameter value is present by itself. Averaging will produce an output corresponding to the average of the two parameter values, and winner-take-all will produce a bimodal distribution of responses as the system chooses between the two potential winners.

Microstimulation is an ideal tool for manipulating the signals in a map in such a way as to distinguish between these possibilities. In one such experiment, Groh, Born, and Newsome used microstimulation to introduce artificial activity into the map of stimulus velocity in area MT of monkeys. At the same time, the monkeys were presented with a moving visual target that they had been trained to track using smooth pursuit eye movements. The visual target was presented inside the positional receptive fields of the cells at the microstimulation site

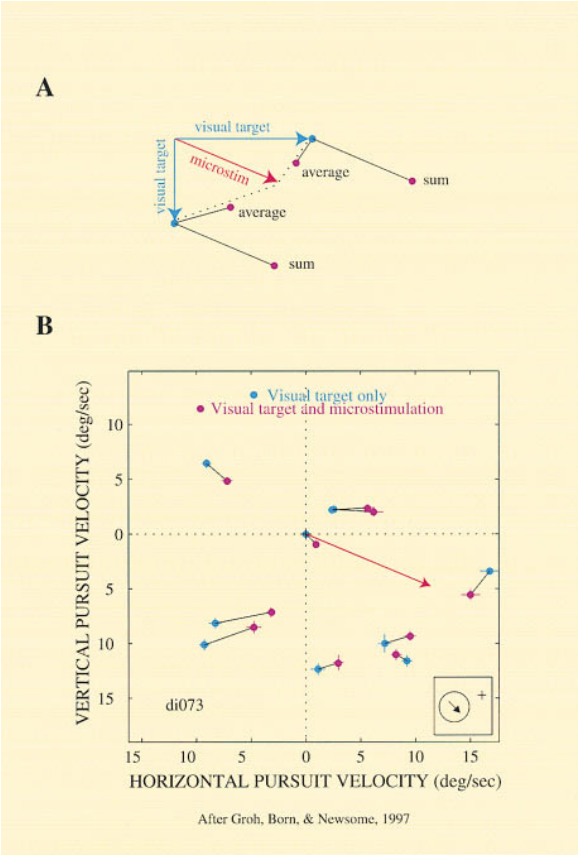

Figure 1. Effects of Microstimulation in Cortical Area MT on Smooth Pursuit Eye Movements

After Groh et al. (1997).

(A) The predicted results of vector averaging and vector summation are shown. Consider a stimulation site in MT where microstimulation introduces a signal corresponding to a downward and rightward velocity (red arrow). When microstimulation is paired with a visual target moving straight down (downward blue arrow) the results (purple) depend on whether a vector average or a vector sum of the visual and electrically induced velocity signals is calculated. If a vector average is calculated, then the smooth pursuit should be shifted to the right and slightly upward on stimulated trials for this target velocity. However, if a vector sum is calculated, smooth pursuit should be shifted down and to the right. Conversely, if the microstimulation is paired with a target moving straight right (rightward blue arrow), the vector average would result in pursuit that was shifted down and slightly to the left, whereas a vector sum would produce the same downward and rightward shift for all target velocities.

(B) Results for a typical stimulation site in area MT. The velocity signal induced by the microstimulation corresponds to a downward and rightward velocity (red arrow). The blue points indicate the mean velocity of pursuit on 15–20 nonstimulated trials for a range of different visual target velocities. Connected to each blue data point is a purple data point that shows the mean pursuit velocity on stimulated trials with the same visual target velocity. The inset shows the location of the positional receptive field and the preferred direction of the cells at the microstimulation site. The overall pattern of results is more consistent with vector averaging—microstimulation shifts the smooth pursuit toward the stimulation-induced velocity signal. The relative change in pursuit varies with the velocity of the target, rather than being consistently shifted down and to the right as would be predicted by vector summation.

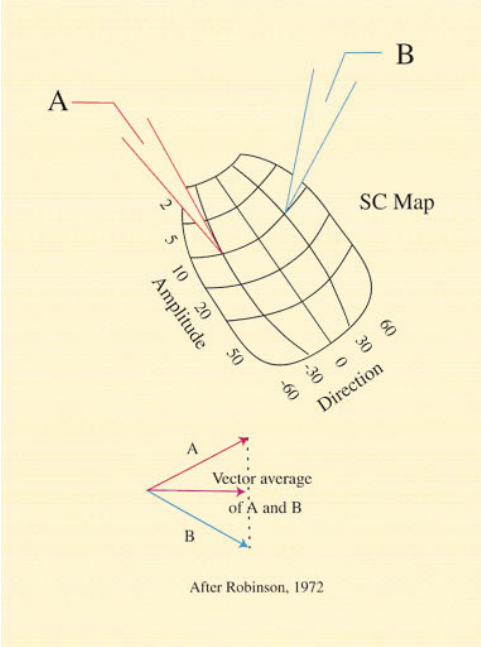

Figure 2. Schematic Diagram of the Results of Microstimulation at Two Locations in the Primate SC

After Robinson (1972). Simultanous stimulation at two sites produces the vector average, not the vector sum, of the saccades evoked by stimulation at either site alone.

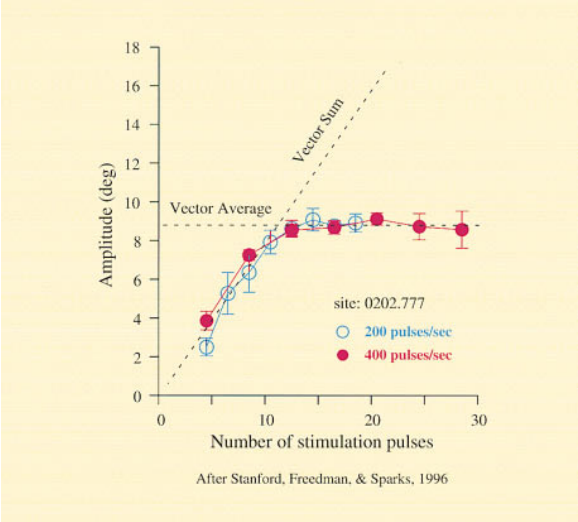

Figure 3. Relationship between the Amplitude of Saccades Evoked by Microstimulation in the Primate SC and the Number of Pulses in the Stimulation Train

After Stanford, Freedman, and Sparks (1996). Data are from one stimulation site in the SC and two different stimulation frequencies. A read-out mechanism that calculated a true vector average would produce the same saccade amplitude regardless of the number of stimulation pulses, while a true vector sum would produce a saccade amplitude that scaled linearly with the number of pulses (dashed lines). This experiment shows that the outcome represents a mixture of these two hypotheses.

in MT. By examining the animal's smooth pursuit in response to this combined visual velocity signal and electrically induced velocity signal, we were able to establish that the smooth pursuit system calculates a vector average of the velocity signals present in area MT (Figure 1).

Related studies in the SC have also suggested that a vector average is calculated when this structure is read out: microstimulation at two sites in the SC elicits a saccade that is the vector average of the saccade elicited by stimulation at either site alone (Figure 2; Robinson, 1972). Similar averaging occurs in response to microstimulation paired with a visual target. Microstimulation in the frontal eye fields (FEF), a structure that encodes saccade vectors in a manner very similar to the SC and shares a similar projection pattern, supports averaging as the read-out for this area as well (see references in Groh et al., 1997). In addition, strictly behavioral experiments have shown that presentation of two saccade targets or two pursuit targets can cause a similar vector averaging pattern (Findlay, 1982; Lisberger and Ferrera, 1997). In fact, the emerging picture is that when the behavioral response is a movement that is under reasonably direct control by the signals in the map, the read-out mechanism usually calculates an average. A few experiments that used different methodology have found evidence for winner-take-all (Salzman and Newsome, 1994; Ferrera and Lisberger, 1995), but evidence for vector summation is quite rare.

Why does the read-out mechanism generally calculate an average rather than a sum? Averages and sums differ only by a scaling factor related to the number of parameter values or the overall level of activity in the map. However, calculating an average is much more sensible, because an average normalizes for fluctuations in activity level or in the number of active neurons. Failure to normalize can have substantial consequences: if two saccade targets were presented right next to one another (either real saccade targets or a stimulation-induced saccade signal), a summation algorithm would produce a saccade of twice the appropriate amplitude.

Thus, calculating an average is clearly a sensible algorithm for the brain to use in reading its internal representations. But how does the brain actually compute the average locus of activity in a map? Insight into this question comes from more detailed studies of the effects of microstimulation in the SC in the barn owl, cat, and monkey (du Lac and Knudsen, 1990; van Opstal et al., 1990; Pare et al., 1994; Stanford et al., 1996). These studies examined the effects of varying the microstimulation frequency, current level, and/or train duration on saccade amplitude and velocity. Varying the stimulation parameters in this manner is another means for distinguishing between vector averaging and vector summation. Just as a vector average mechanism predicts that stimulation at two sites should produce an output that corresponds to the average of the two, stimulation at one site should produce an output corresponding to the average location of the activated cells at that single site. Varying the current level should change the size of the active population but not the average location of that activity. Varying the stimulation frequency should change the discharge rate of the stimulated neurons, but again, it should have no effect on the average location of activity. Stimulation duration should affect only the duration of

stimulation-induced activity but not its location. In contrast, if the activity in the SC were summed by a vector summation mechanism, the output should scale with all of these parameters.

The results of these experiments partly confirmed the results of two-site stimulation experiments: above a certain frequency and current level, the output, namely the amplitude of the saccade, was constant. Additional increases in these parameters did not increase movement amplitude. However, reducing the frequency, current level, or duration shortened the amplitude of the movement. Stimulation frequency also influenced the latency and velocity of the movement—higher frequency stimulation triggered the movement sooner and produced a faster movement, while lower frequency stimulation produced a later, slower movement. Stanford et al. (1996) further demonstrated a trade-off between frequency and duration, showing that within broad limits, saccade amplitude scales with the number of pulses (the product of frequency and duration) until the "site-specific amplitude" is reached (Figure 3).

These results provide fascinating clues to the actual mechanism for reading SC. Models that seek to explain the read-out of the SC must simultaneously produce both the dependence of the output on stimulation parameters below a certain threshold and the independence or normalization above that level. Given the overall similarities between the SC, FEF, and MT, as well as numerous other systems, it will be interesting to explore whether a similar pattern holds true in these areas as well.

On a final note, the very fact that the signals produced by microstimulation in the brain are artificial and might never occur under normal conditions is part of what makes this technique so informative. Microstimulation allows the neurophysiologist to explore how the system responds to unusual perturbations. Just as sensory illusions trick the brain into revealing its mechanisms, abnormal neural activity patterns can reveal how the system is designed to function normally. In concert with recording and lesion experiments, microstimulation is a critical, and highly informative, technique for exploring neural representations and how they are read out.

Selected Reading

Carpenter, R.H.S. (1988). Movements of the Eyes (London: Pion).

du Lac, S., and Knudsen, E.I. (1990). J. Neurophysiol. *63*, 131–146.

Ferrera, V.P., and Lisberger, S.G. (1995). J. Neurosci. *15*, 7472–7484.

Findlay, J.M. (1982). Vision Res. *22*, 1033–1045.

Groh, J.M., Born, R.T., and Newsome, W.T. (1997). J. Neurosci. *17*, 4312–4330.

Lisberger, S., and Ferrera, V. (1997). J. Neurosci. *17*, 7490–7502.

Pare, M., Crommelinck, M., and Guitton, D. (1994). Exp. Brain Res. *101*, 123–139.

Robinson, D.A. (1972). Vision Res. *12*, 1795–1807.

Salzman, C.D., and Newsome, W.T. (1994). Science *264*, 231–237.

Sparks, D.L., and Hartwich-Young, R. (1989). In The Neurobiology of Saccadic Eye Movements, R.H. Wurtz and M.E. Goldberg, eds. (New York: Elsevier), pp. 213–255.

Stanford, T.R., Freedman, E.G., and Sparks, D.L. (1996). J. Neurophysiol. *76*, 3360–3381.

van Opstal, A.J., van Gisbergen, J.A.M., and Smit, A.C. (1990). Exp. Brain Res. *79*, 299–312.