

Fast Remapping of Sensory Stimuli onto Motor Actions on the Basis of Contextual Modulation

Emilio Salinas

Department of Neurobiology and Anatomy, Wake Forest University School of Medicine, Winston-Salem, North Carolina 27157-1010

Higher organisms can establish complex associations between sensory events and motor responses. More remarkable than their complexity, however, is that the resulting sensory-motor maps can be selectively interchanged. For example, a person who speaks English and Spanish can read aloud “con once, sin once,” going effortlessly from one language to the other. What is the neural basis of this capacity? Here, a network model is presented in which multiple maps between sensory stimuli and motor actions are possible, but only one of them, depending on behavioral context, is implemented at any given time. The key is a nonlinear representation in which the gain of sensory responses is regulated by context information. Neuronal responses can indeed show variations in gain, as has been documented in the case of proprioceptive signals such as eye and head position, which can modulate visually triggered activity. However, in contrast to these, the contextual cues used here need not bear any relationship to the physical attributes of the stimuli; in particular, spatial location is irrelevant. The model thus postulates the existence of sensory neurons that are nonlinearly modulated by arbitrary context signals, a plausible and testable prediction. The proposed mechanism allows a network of neurons to effectively change the functional connectivity between its inputs and outputs and may partially explain how animals can quickly adapt their behavior to varying environmental conditions.

Key words: sensory-motor integration; gain modulation; coordinate transformation; arbitrary visuomotor remapping; neural coding; basis functions; neural network

Introduction

If somebody suddenly holds your hand, your reaction may vary dramatically depending on whether the person who touches you is a nurse taking your pulse, a son grabbing your attention, or a stranger with unknown intentions. Clearly, the impact of a perceived stimulus depends on other concurrent stimuli, recent and distant events, and current motivations and goals (Drea, 1998; Platt and Glimcher, 1999; Handel and Glimcher, 2000; Wise and Murray, 2000; Hobin et al., 2003). How is all of this context information combined and used to influence motor behavior?

A generic way by which the nervous system integrates information from various modalities or sources is through changes in gain (Andersen et al., 1997; Salinas and Thier, 2000; Salinas and Abbott, 2001; Salinas and Sejnowski, 2001). Neurons with so-called “gain fields” react to certain sensory stimuli, but their overall responsiveness varies as a function of some modulatory parameter, which does not affect their selectivity. Many cases of gain-modulated (GM) responses have been documented, and theoretical work has shown that, in general, they are useful for performing coordinate transformations (Zipser and Andersen, 1988; Andersen et al., 1993; Salinas and Abbott, 1995, 1997;

Pouget and Sejnowski, 1997; Deneve and Pouget, 2003). For instance, eye-centered visual receptive fields that are gain modulated by the direction of gaze can give rise to body-centered receptive fields downstream (Zipser and Andersen, 1988; Salinas and Abbott, 1995, 1997; Pouget and Sejnowski, 1997). However, previous models and neurophysiological experiments have focused almost exclusively on modulation by internal or proprioceptive signals, such as gaze direction (Andersen and Mountcastle, 1983; Andersen et al., 1985, 1990; Brotchie et al., 1995), eye or head velocity (Snyder et al., 1998; Shenoy et al., 1999), arm position (Buneo et al., 2002), or attentional location (Connor et al., 1997; McAdams and Maunsell, 1999; Treue and Martínez-Trujillo, 1999), which are directly combined with spatial sensory information to produce a change in reference frame (Andersen et al., 1993; Pouget and Snyder, 2000; Salinas and Thier, 2000; Salinas and Abbott, 2001).

This study suggests that the nervous system might generalize this strategy to situations in which the transformation is more abstract and depends on arbitrary context signals that contain no intrinsic spatial value. For example, consider a musician who can play the flute and the violin; the stimulus is the score of a tune. While playing, the printed notes are translated into finger and mouth movements that are specific for each instrument. In switching instruments, the functional connectivity between visual and motor networks in the brain changes drastically and almost instantaneously. The contextual input does not shape the sensory-motor map, however; it simply identifies the correct one. That is, context indicates which instrument to play but not how to play it.

To establish a clear link with possible neurophysiological ex-

Received Oct. 9, 2003; revised Nov. 24, 2003; accepted Nov. 25, 2003.

This research was supported by startup funds from Wake Forest University School of Medicine and National Institute of Neurological Disorders and Stroke Grant NS044894-01 to E.S. I thank Nick Bentley for stimulating discussions and Terry Stanford and Alex Pouget for their suggestions and critical comments.

Correspondence should be addressed to Emilio Salinas, Department of Neurobiology and Anatomy, Wake Forest University School of Medicine, Winston-Salem, NC 27157-1010. E-mail: esalinas@wfbmc.edu.

DOI:10.1523/JNEUROSCI.4569-03.2004

Copyright © 2004 Society for Neuroscience 0270-6474/04/241113-06\$15.00/0

periments, the results are presented here in terms of a hypothetical remapping task that requires the same type of sensory-motor flexibility just discussed but is well suited for experimental work with awake primates. The computer simulations presented here show that a network with sensory neurons that are gain modulated by context cues can solve this task efficiently.

Materials and Methods

All simulations were performed using Matlab (The Mathworks, Natick, MA). The code is available on request.

Visuomotor task. The task used to investigate the remapping problem is schematized in Figure 1. In any trial of this task, 1 of 16 stimuli (Fig. 1e) is presented, and the subject has to classify it. There are four possible response targets, and the chosen class is indicated by an eye movement made toward one of the targets. Crucially, there are five task conditions, four that generate unique maps between stimuli and responses plus an additional no-go condition in which a stimulus is presented but no movement is made. The condition that applies in each trial is indicated by a separate nonspatial cue, the color of the fixation spot, which acts as the context.

Network architecture. The network that simulates the task has two layers. The first contains GM neurons, and these drive a second layer of output or motor neurons through a set of connections. Each GM unit makes contact with all output units. In each trial of the task, the GM neurons are activated by the sensory and context signals, and a movement is generated by the output neurons. For this, the profile of activity of the output neurons should have a single peak indicating the location (or direction) of the movement to be made. Note that the temporal development of the task is not modeled, just the sensory-motor associations.

GM responses. The GM neurons respond to the stimuli, but their selectivity for individual features, including length, color, and orientation, is not modeled explicitly. A particular stimulus is simply a label between 1 and 16 that evokes a certain amount of activity in each GM unit. Although actual neuronal responses typically vary smoothly as functions of single sensory features, this is a reasonable simplification when only a few stimuli are selected from a vast, multidimensional feature space, as in the present task.

The stimulus is indicated by x , which takes integer values from 1 to 16. The context or condition is indicated by y , which takes integer values from 1 to 5. The mean firing rate, r_j , of GM unit j depends on x and y , and is written as follows:

$$r_j = r_{\max} f_j(x) [1 - D + D g_j(y)] + B, \quad (1)$$

where $f_j(x)$ and $g_j(y)$ are functions that vary between 0 and 1, B is a baseline rate equal to 4 spikes/sec, $r_{\max} = 35$ spikes/sec, and D is the modulation depth. Throughout this study, $D = 0.5$, which produces a maximum contextual suppression of 50%. To include neuronal variability, Gaussian noise is added independently to each GM response in every trial of the task. The noise is multiplicative: the variance of the noise for unit j is equal to αr_j (see Fig. 4 legend).

The gain functions $g_j(y)$ take five values, one for each task condition. These values are $g_j(y) = \{1, 0.8, 0.5, 0.3, 0\}$. Crucially, they are assigned randomly to each of the five conditions, with a new random permutation for each GM unit. As a final step, the $g_j(y)$ values are jittered by small random amounts. In this way, context affects each GM neuron differently. The tuning functions $f_j(x)$ are generated using the same procedure but different numbers. First, 16 preset values between 0 and 1 are generated. Second, for each unit j , these are dealt randomly to the 16 stimuli, with a different permutation for each GM neuron. Finally, random jitter is added to each $f_j(x)$. This guarantees that there is no topography, and that all neurons have different tuning functions.

Output responses. The firing rate R_i of output unit i is calculated through a weighted sum of GM rates:

$$R_i = \sum_j w_{ij} r_j, \quad (2)$$

where w_{ij} represents the synaptic connection from GM neuron j to output neuron i . This expression is used when the GM neurons drive the output neurons. However, there is also an intended or desired response for each output neuron, $F_i(x, y)$, which is used only when setting the connections:

$$F_i(x, y) = \begin{cases} r_{\max} \exp\left(-\frac{[T(x, y) - c_i]^2}{2\sigma^2}\right) + B & \text{if } y \neq 5 \\ B & \text{if } y = 5 \end{cases}, \quad (3)$$

where $y = 5$ is the no-go condition, $\sigma = 0.35$ and c_i , the preferred target location of unit i , varies uniformly between -3 and 3 . The intended output profile of activity is obtained by plotting the above expression as a function of c_i . In go trials, this corresponds to a Gaussian profile that peaks at $T(x, y)$, which is the target location ($-2, -1, +1, \text{ or } +2$) that corresponds to the correct classification of stimulus x in condition y (see Fig. 1). In no-go trials, all output neurons should keep firing at the baseline rate B . The number of output units is always 30. The center of mass of the output activity T_{out} is interpreted as the encoded target location, and is given by:

$$T_{\text{out}} = \frac{\sum_i (R_i - B)^2 c_i}{\sum_k (R_k - B)^2}. \quad (4)$$

The error is equal to $T(x, y) - T_{\text{out}}$, and is computed in go trials only.

Synaptic weights. The connections to any given output neuron are chosen so that they minimize the average squared difference between intended and driven responses. To simplify the notation, consider only one output neuron with driven response $\sum_j w_j r_j$ and intended response $F(x, y)$. The row vector \vec{w} contains all of the connections to this chosen postsynaptic neuron, and the optimal values are given by:

$$\vec{w} = \vec{L} C^{-1}, \quad (5)$$

where

$$C_{jk} \equiv \langle r_j r_k \rangle \quad (6)$$

$$L_k \equiv \langle F(x, y) r_k \rangle. \quad (7)$$

Here, angle brackets indicate an average over all values of x and y and over multiple trials, and C^{-1} is the inverse of the correlation matrix C . This inverse (or the pseudoinverse) is found numerically. To compute the above averages, all GM rates are evaluated using Equation 1 plus the noise terms. Because the noise is multiplicative and uncorrelated across neurons, it contributes an amount $\alpha \langle r_j \rangle$ to each diagonal element C_{ii} . The connections to other output units are computed in the same way, with different \vec{w} vectors but the same matrix C^{-1} .

Additional modifications. The basic model just described is also tested with four additional manipulations. The first is to make the tuning and gain functions binary, which restricts the mean GM responses to three values. That is, the five gain factors for neuron j are now $g_j(y) = \{1, 1, 1, 0, 0\}$, again with a random permutation for each neuron. Similarly, $f_j(x)$ consists of n ones and $16 - n$ zeros, randomly permuted. Under these conditions, Equation 1 can generate only three rates: $r_{\max} + B$, $r_{\max}(1 - D) + B$, and B . Performance in the simulations depends on n , which determines the frequencies of the three rates.

The second manipulation is to use correlated instead of independent noise. The response of GM unit j is equal to $r_j + \sigma_j \gamma_j$, which is the sum of the mean rate plus a noise term; σ_j is the SD of the response, and γ_j is drawn from a Gaussian distribution with zero mean and unit variance, that is, $\langle \gamma_j \rangle = 0$ and $\langle \gamma_j^2 \rangle = 1$. The key to generating correlated samples is to determine the matrix of desired correlation coefficients, which has entries $\langle \gamma_j \gamma_k \rangle$. All elements along the diagonal are equal to 1. Off the diagonal, when samples are drawn independently, $\langle \gamma_j \gamma_k \rangle = 0$. When all pairs are equally correlated, $\langle \gamma_j \gamma_k \rangle = \rho$, and when correlations are proportional to the overlap between response curves:

$$\langle \gamma_j \gamma_k \rangle = \rho \frac{\langle (r_j - \langle r_j \rangle)(r_k - \langle r_k \rangle) \rangle}{\sqrt{\langle (r_j - \langle r_j \rangle)^2 \rangle \langle (r_k - \langle r_k \rangle)^2 \rangle}} \quad (8)$$




Figure 1. Visuomotor remapping task. The sequence of events in each trial is as follows: a stimulus (bar) is presented while the subject stares at a fixation point (colored dot); the stimulus disappears; four targets appear (gray dots), and the subject classifies the stimulus, indicating his choice with an eye movement (horizontal arrow) toward one of the targets. The stimulus feature being classified and the correspondence with the targets are cued by the color of the fixation spot. *a*, In condition 1, the fixation spot is magenta, the classification criterion is orientation, and left and right proximal targets are used for horizontal and vertical bars, respectively. *b*, In condition 2, the fixation spot is orange, and stimuli are classified by orientation, as in condition 1 (*a*) but with the targets reversed. *c*, In condition 3, the fixation spot is cyan, the classification criterion is color, and left and right eccentric targets are for red and blue bars, respectively. *d*, In condition 4, the fixation spot is green, and stimuli are classified by color, as in condition 3 (*c*), but with the targets reversed. Condition 5 is no-go (data not shown; yellow spot); a stimulus is presented but no movement is made. Arrows indicate correct movements. *e*, Stimuli, numbered 1–16, that may be used in the task. The model requires 16 items that, like these, are distinct but has no knowledge about their individual features (see Materials and Methods).

where ρ measures the correlation strength and is between 0 and 1. Given this matrix, correlated samples can be obtained, for instance, by calling the Matlab function `mvrnd`. Although noise correlations can also be included in the calculation of synaptic weights, this makes little difference; thus, results are reported using the unmodified method described above.

In the third manipulation, the GM cells combine sensory and context signals linearly. In this case, the mean firing rate of GM cell j is given by:

$$r_j = \frac{r_{\max}}{2} [f_j(x) + g_j(y)] + B, \quad (9)$$

instead of Equation 1.

Finally, an alternative nonlinear interaction between stimulus and context is implemented by substituting Equation 1 with:

$$r_j = r_{\max}((1 - D)f_j(x) + D[f_j(x) + g_j(y) - 1]_+) + B, \quad (10)$$

where the brackets indicate rectification; that is, $[x]_+ = \max\{0, x\}$. In each case, everything else is as in the original simulations.

Results

The model network performs the visuomotor task shown in Figure 1. In each trial, the GM neurons respond to the sensory and context signals and drive a set of output or motor neurons, which then generate an eye movement.

The responses of two representative GM units are displayed in Figure 2, *a* and *c*. The plots show the firing rates evoked by each stimulus, with different colors corresponding to different task conditions or contexts. The cells respond more strongly to some stimuli than to others, and their preference order is different. In Figure 2*a*, the three most effective stimuli are 9, 14, and 10, whereas in Figure 2*b*, they are 3, 15, and 2. In the network, the preference sequences are set randomly, so each cell ends up with a random stimulus–response curve, or tuning curve, that is unique. However, note that these tuning curves have the same shape for the five modulatory conditions. This is crucial and reflects the key assumption that context influences the sensory responses by changing their gain; thus, the five curves differ only in their amplitudes. Indeed, this is achieved by modeling each GM response as a product of two factors, one that depends only on the stimulus and another that depends only on the condition (see Materials and Methods). In this way, context affects the over-

all responsiveness of the neuron but not its selectivity, which is the defining feature of gain modulation (Salinas and Thier, 2000; Salinas and Abbott, 2001; Salinas and Sejnowski, 2001). Plotting the GM responses in the format of Figure 2, *b* and *d*, is an alternative way to reveal gain variations (Fig. 2 legend) (McAdams and Maunsell, 1999). Another important feature of the GM units is that the order in which they prefer the five conditions is set randomly, so it is also different for each cell (see Materials and Methods).

The firing rate of each output neuron in the model is determined by a weighted sum of GM rates in which the weights represent synaptic connections. The output population is meant to encode the location of a target for an impending movement. Thus, in every go trial, the evoked profile of activity should have a single peak indicating the correct target location, which depends on the stimulus and the context of

that trial. In the no-go condition, however, all output responses should stay at their low baseline level, so the profile should be flat. The connections that achieve all of this as best as possible are found through an optimal algorithm (see Materials and Methods). This algorithm is run only once; afterward, the synaptic weights are not adjusted any further.

Having specified the GM tuning curves and the network connections, the model is tested in a series of trials of the task. Each trial proceeds by (1) choosing a stimulus and a context, (2) generating all GM responses (Eq. 1), (3) calculating the driven, output firing rates (Eq. 2), and (4) calculating the center of mass of the output activity (Eq. 4), which is taken as the encoded target location. Finally, the encoded location is compared with the location that should have been reached given the stimulus and the condition. It is important to note that noise is added to all GM responses in each trial of the task (see Materials and Methods). This is to simulate neuronal variability and evaluate the robustness of the model.

Figure 3 illustrates the behavior of the full model in four single trials. Here, 30 output neurons are driven by 864 GM units. The GM responses are color coded and, for graphical purposes, are ordered according to preferred stimuli. There is no single focus of intense activity, only a diffuse band. This is because the neurons preferring, for example, stimulus 8 are clustered together, but the neurons that have stimulus 8 as their second or third preferred stimulus are randomly scattered. The output neurons, ordered by their preferred target locations, are meant to produce a peak of activity at either of four points, -2 , -1 , $+1$, or $+2$, which correspond to the four response targets in the task (Fig. 1). In some trials, there is a clear single peak of activity at or very near the proper target location (Fig. 3*a,b*). In other trials, secondary peaks are evident (Fig. 3*c*). These arise because of the noise added to the GM responses; they entirely disappear with zero noise. Their effect is to shift the center of mass, increasing the error. Overall, however, the model performs accurately, because the encoded target location (red vertical line) is typically very close to the intended one (black vertical line) for all combinations of stimulus and condition.

The model also performs correctly in the no-go condition; the profile of output responses is approximately flat, with rates close

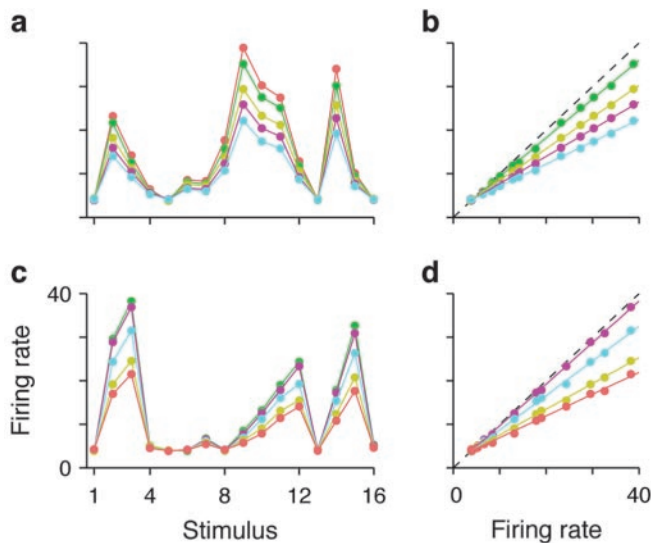


Figure 2. Responses of model GM neurons. *a*, Mean firing rates evoked by each of the 16 stimuli for one GM unit. Each color corresponds to a task condition. The preferred condition of this neuron is orange. *b*, Same data as in *a*, but plotting the firing rates in the preferred condition on the *x*-axis and the firing rates in the nonpreferred conditions for the corresponding stimuli on the *y*-axis. Because the rates in preferred and nonpreferred conditions are proportional, the resulting points lie along straight lines. *c*, *d*, As in *a* and *b*, but for a different model neuron, the gain of which is maximal in the green condition. Firing rate units are spikes per second.




Figure 3. Network performance in single trials. Left column shows all 864 GM responses, color coded and sorted (for display purposes only) according to their most effective stimuli. Right column shows all 30 output responses arranged by their preferred target locations. *a*, Trial with stimulus 4 in condition 1. Black and red vertical lines indicate intended and encoded target locations, respectively; error is -0.05 (arbitrary distance units). *b*, Trial with stimulus 4 in condition 2; error is 0.02 . *c*, Trial with stimulus 12 in condition 3; error is -0.36 . *d*, Trial with stimulus 12 in condition 5 (no-go). The rms error averaged over all stimuli and go conditions was 0.22 . Response units are spikes per second.

to the baseline firing level (Fig. 3*d*). This can be quantified by measuring the highest firing rate of all motor neurons in go versus no-go trials. For the network of 864 GM neurons, in no-go trials the maximum rate was on average 8.9 ± 2.5 (SD) spikes/sec, whereas in go trials, it was 35.6 ± 4.2 spikes/sec. Also, without noise, the profile of output responses was perfectly flat.

The accuracy of the network in go trials depends on the variability of the GM neurons. That is, the difference between correct and encoded target locations (the error) increases with noise. In the simulations in Figure 3, the variance of the noise of each unit was equal to the mean response of the unit, and the ensuing root mean square (rms) error was 0.2 (the average error was practically 0, because positive and negative deviations were equally likely); for reference, recall that the maximum separation between targets is 4 units. One way to measure the robustness of the model with respect to noise is to plot the rms error as a function of the number of GM neurons, N , as in Figure 4*a*. The error decreases sharply (faster than $1/\sqrt{N}$) indicating that noise is what limits the performance of the network. In accordance, when no noise was added, the error was virtually zero. The performance of the network can also be quantified by considering a trial successful whenever the encoded target location is within a certain fixed distance from the proper target. In this way, a percentage of incorrect classifications or incorrect movements can be measured. This classification error also drops sharply with network size (Fig. 4*b*).

Discussion

The model presented here achieves drastic, context-dependent changes in functional connectivity, which is a way of folding multiple networks, each rendering a specific sensory-motor map, into one. The underlying mechanism is the same as that proposed for performing coordinate transformations from one reference frame to another (Zipser and Andersen, 1988; Andersen et al., 1993; Salinas and Abbott, 1995, 1997; Pouget and Sejnowski, 1997; Deneve and Pouget, 2003) and has the great advantage that the output neurons can read out the correct maps using a simple algorithm, a weighted sum. However, three novel results stand out. First, previous modeling studies typically aimed to combine a modulatory quantity with some feature of a stimulus. The classic example is adding eye position to stimulus position in retinal coordinates to obtain the location of the stimulus in head-centered coordinates (Zipser and Andersen, 1988; Salinas and Abbott, 1995; Pouget and Sejnowski, 1997). In contrast, here, the stimuli and modulatory cues do not need to be related or combined in any particular way. Second, no sensory topography is required. Variants of the model in which the GM responses change smoothly across the population or across stimuli also work, and this smoothness is important for generalization, that is, for correctly classifying novel, unseen stimuli (data not shown). However, it is not strictly necessary for selecting among previously learned maps, as this model does. Finally, the dramatic changes in motor activity can be driven by relatively subtle variations in the firing of the GM cells: the maximum decrease in a GM rate between two conditions was 50% (see Materials and Methods), and the average decrease was $\sim 25\%$ (Fig. 3, compare *a* and *b*). These numbers are well inside the range of experimental values reported for modulatory variables such as attention (Connor et al., 1997; McAdams and Maunsell, 1999; Treue and Martínez-Trujillo, 1999).

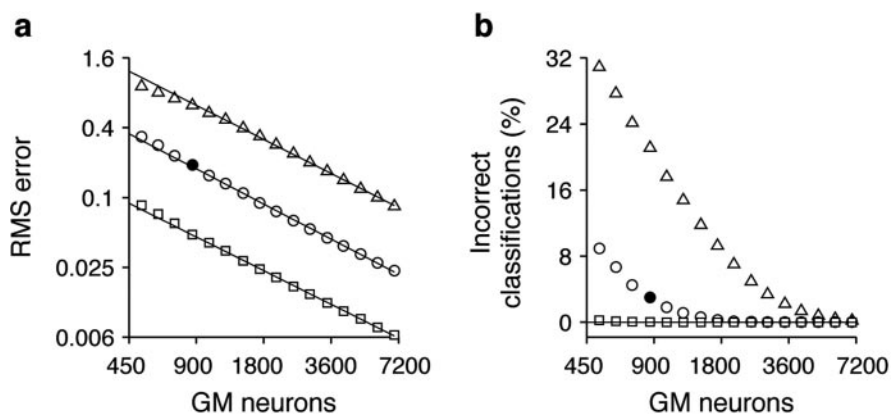


Figure 4. Task performance as a function of network size and response variability. *a*, rms error versus the number of GM neurons in the network; note log scales. For each GM unit, the variance of the noise was equal to α times the mean response (see Materials and Methods). The three curves correspond to different noise levels: $\alpha = 0.25$ (squares); $\alpha = 1$ (circles); and $\alpha = 4$ (triangles). Noise was uncorrelated across neurons. Straight lines are fits to the data above 800 GM neurons; slopes are approximately -1 . *b*, Percentage of incorrect classifications versus the number of GM neurons, for the same simulations in *a* (note log-linear scales). A trial was deemed incorrect if the encoded target location was >0.5 U away from the correct target. Filled symbols correspond to the network used in Figure 3.

Robustness and generality of the model

The behavior of the model is highly insensitive to many details of its implementation; other parameter choices, tuning curves, and gain functions produce essentially the same results. For example, instead of using the full range from 4 to 39 spikes/sec as in Figure 2*a*, the average GM responses can be restricted to only three possible rates, 4, 21.5, and 39 spikes/sec (see Materials and Methods). This apparently harsh manipulation may shift the curves in Figure 4*a* up or down by a factor of 2, depending on the relative frequencies of those rates, but it preserves the inverse relationship between error and network size (with the same slopes). Therefore, the qualitative behavior of the network is unchanged.

One aspect of the model that could be refined for achieving higher accuracy is the output layer. In particular, recurrent connections between output units could be used to reinforce and smooth the unimodal profile of activity, eliminating secondary bumps like those in Figure 3, *a–c* (Salinas and Abbott, 1996; Deneve et al., 1999). This could be achieved with a simple center-surround organization and would make the readout performed by this layer optimal (Deneve et al., 1999). Although significant, the difference would again be quantitative.

However, there is one factor that can potentially generate a minimum rms error that cannot be eliminated: the presence of noise correlations (Zohary et al., 1994). These arise when the firing rates of pairs of neurons tend to fluctuate together across identical trials. Their impact depends not only on their strength and distribution across a population but also on how the neuronal responses are combined postsynaptically (Abbott and Dayan, 1999; Romo et al., 2003). In the best-case scenario, the common noise is cancelled, producing an increase in performance. In the worst-case scenario, part of the common variability remains, regardless of network size, and thus a hard limit is imposed. In the model, introducing a constant correlation coefficient of 0.15 (see Materials and Methods) (Zohary et al., 1994; Romo et al., 2003) between all pairs of GM neurons resulted in slightly smaller rms errors. This is because each postsynaptic neuron combines many GM rates with positive and negative weights, so the fluctuations tend to average out. Stronger correlations and mixtures of positive and negative values led to similar results. Smaller errors were observed even when each correlation coefficient was propor-

tional to the overlap between response curves (Eq. 8), a more plausible situation. In a real biological circuit, correlations could still conspire to limit the accuracy of the maps, but the results suggest that the distributed nature of the network makes it rather resistant to this problem.

The basic two-layer model is quite general, in that it works with many other remapping tasks (data not shown). The key is the joint sensory–contextual representation; the GM neurons act as basis functions from which other, arbitrary functions of stimulus and context can be constructed (Poggio, 1990; Pouget and Sejnowski, 1997; Pouget and Snyder, 2000; Ben Hamed et al., 2003). There are two important aspects to this: the GM array must span all relevant combinations of sensory stimuli and modulatory signals, and the joint encoding of sensory and modulatory influences must be nonlinear (Salinas and Abbott, 1995, 1997; Pouget and Sejnowski, 1997; Deneve and Pouget, 2003). This last condition was verified in two ways. When the multiplication of sensory- and context-dependent terms that determines the GM firing rates was substituted with a simple addition (Eq. 9), all mappings failed miserably. For the network in Figure 3, the rms error went from 0.22 to 1.6, and the classification error increased from 3 to 94%. In contrast, when the multiplication was substituted with a rectification of the sum (Eq. 10), the rms and classification errors decreased slightly, to 0.19 and 1.5%, respectively. Therefore, a perfect multiplication is not necessary. Interestingly, however, experiments have reported neuronal interactions that seem to approximate an exact multiplication (Brotchie et al., 1995; McAdams and Maunsell, 1999; Treue and Martínez-Trujillo, 1999; Peña and Konishi, 2001). This may be useful for learning the connections between GM and motor units using simple Hebbian mechanisms (Salinas and Abbott, 1995, 1997), so the optimal nonlinearity may also depend on the available procedures for synaptic modification.

Experimental predictions

Various neurophysiological studies have explored changes in neuronal properties as functions of multiple sensory cues in the spirit of the task proposed here (White and Wise, 1999; Lauwereyns et al., 2001; Wallis and Miller, 2003). Some of those results are consistent with context-dependent variations in gain but, because the paradigms were not designed to assess this, confounding factors are inevitable. To test the network model, the present task was somewhat more complex than in actual primate experiments, but it may be simplified considerably by using fewer stimuli and conditions as long as plots like those of Figure 2 are still possible. Regardless of the specifics of the behavioral task, the cleanest demonstration of the proposed mechanism would be a pure gain effect, which would show up as a context-dependent variation in firing intensity that leaves stimulus selectivity intact, as in Figure 2.

However, other nonlinear interactions are certainly possible, as documented with sensory responses that are modulated by eye position, hand position, or eye velocity (Jay and Sparks, 1987; Buneo et al., 2002; Ben Hamed et al., 2003). The most general prediction of the remapping model is that neuronal responses in

one or more cortical areas should behave as basis functions (Ben Hamed et al., 2003), meaning that (1) most of them should mix sensory and context signals, (2) the mixing should be nonlinear, and (3) a linear combination of responses (Eq. 2) should perform extremely well at reading out an arbitrary function of stimulus and context. Whether purely multiplicative or not, such nonlinear responses would extend the generality of gain modulation and basis–function expansion as a computational strategy in the brain (Poggio, 1990; Salinas and Thier, 2000; Pouget and Snyder, 2000) and would partly explain the great flexibility with which higher organisms adapt their actions to current environmental conditions.

References

- Abbott LF, Dayan P (1999) The effect of correlated activity on the accuracy of a population code. *Neural Comput* 11:91–101.
- Andersen RA, Mountcastle VB (1983) The influence of the angle of gaze upon the excitability of light-sensitive neurons of the posterior parietal cortex. *J Neurosci* 3:532–548.
- Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. *Science* 230:450–458.
- Andersen RA, Bracewell RM, Barash S, Gnadt JW, Fogassi L (1990) Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J Neurosci* 10:1176–1198.
- Andersen RA, Snyder LH, Li C-S, Striccanne B (1993) Coordinate transformations in the representation of spatial information. *Curr Opin Biol* 3:171–176.
- Andersen RA, Snyder LH, Bradley DC, Xing J (1997) Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annu Rev Neurosci* 20:303–330.
- Ben Hamed S, Page W, Duffy C, Pouget A (2003) MSTd neuronal basis functions for the population encoding of heading direction. *J Neurophysiol* 90:549–558.
- Brotchie PR, Andersen RA, Snyder LH, Goodman SJ (1995) Head position signals used by parietal neurons to encode locations of visual stimuli. *Nature* 375:232–235.
- Buneo CA, Jarvis MR, Batista AP, Andersen RA (2002) Direct visuomotor transformations for reaching. *Nature* 416:632–636.
- Connor CE, Preddie DC, Gallant JL, Van Essen DC (1997) Spatial attention effects in macaque area V4. *J Neurosci* 17:3201–3214.
- Deneve S, Pouget A (2003) Basis functions for object-centered representations. *Neuron* 37:347–359.
- Deneve S, Latham PE, Pouget A (1999) Reading population codes: a neural implementation of ideal observers. *Nat Neurosci* 2:740–745.
- Drea CM (1998) Social context affects how rhesus monkeys explore their environment. *Am J Primatol* 44:205–214.
- Handel A, Glimcher PW (2000) Contextual modulation of substantia nigra pars reticulata neurons. *J Neurophysiol* 83:3042–3048.
- Hobin JA, Goosens KA, Maren S (2003) Context-dependent neuronal activity in the lateral amygdala represents fear memories after extinction. *J Neurosci* 23:8410–8416.
- Jay MF, Sparks DL (1987) Sensorimotor integration in the primate superior colliculus. II. Coordinates of auditory signals. *J Neurophysiol* 57:35–55.
- Lauwereyns J, Sakagami M, Tsutsui K, Kobayashi S, Koizumi M, Hikosaka O (2001) Responses to task-irrelevant visual features by primate prefrontal neurons. *J Neurophysiol* 86:2001–2010.
- McAdams CJ, Maunsell JHR (1999) Effects of attention on orientation tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441.
- Peña JL, Konishi M (2001) Auditory spatial receptive fields created by multiplication. *Science* 292:249–252.
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238.
- Poggio T (1990) A theory of how the brain might work. *Cold Spring Harbor Symp Quant Biol* 55:899–910.
- Pouget A, Sejnowski TJ (1997) Spatial transformations in the parietal cortex using basis functions. *J Cogn Neurosci* 9:222–237.
- Pouget A, Snyder LH (2000) Computational approaches to sensorimotor transformations. *Nat Neurosci* 3[Suppl]:1192–1198.
- Romo R, Hernández A, Zainos A, Salinas E (2003) Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron* 38:649–657.
- Salinas E, Abbott LF (1995) Transfer of coded information from sensory to motor networks. *J Neurosci* 15:6461–6474.
- Salinas E, Abbott LF (1996) A model of multiplicative neural responses in parietal cortex. *Proc Natl Acad Sci USA* 93:11956–11961.
- Salinas E, Abbott LF (1997) Invariant visual responses from attentional gain fields. *J Neurophysiol* 77:3267–3272.
- Salinas E, Abbott LF (2001) Coordinate transformations in the visual system: how to generate gain fields and what to compute with them. *Prog Brain Res* 130:175–190.
- Salinas E, Sejnowski TJ (2001) Gain modulation in the central nervous system: where behavior, neurophysiology and computation meet. *Neuroscientist* 7:430–440.
- Salinas E, Thier P (2000) Gain modulation: a major computational principle of the central nervous system. *Neuron* 27:15–21.
- Shenoy KV, Bradley DC, Andersen RA (1999) Influence of gaze rotation on the visual response of primate MSTd neurons. *J Neurophysiol* 81:2764–2786.
- Snyder LH, Grieve KL, Brotchie P, Andersen RA (1998) Separate body- and world-referenced representations of visual space in parietal cortex. *Nature* 394:887–891.
- Treue S, Martínez-Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579.
- Wallis JD, Miller EK (2003) From rule to response: neuronal processes in the premotor and prefrontal cortex. *J Neurophysiol* 90:1790–1806.
- White IM, Wise SP (1999) Rule-dependent neuronal activity in the prefrontal cortex. *Exp Brain Res* 126:315–335.
- Wise SP, Murray EA (2000) Arbitrary associations between antecedents and actions. *Trends Neurosci* 23:271–276.
- Zipser D, Andersen RA (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679–684.
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143.