

Transfer of Coded Information from Sensory to Motor Networks

Emilio Salinas and L. F. Abbott

Center for Complex Systems, Brandeis University, Waltham, Massachusetts 02254

During sensory-guided motor tasks, information must be transferred from arrays of neurons coding target location to motor networks that generate and control movement. We address two basic questions about this information transfer. First, what mechanisms assure that the different neural representations align properly so that activity in the sensory network representing target location evokes a motor response generating accurate movement toward the target? Coordinate transformations may be needed to put the sensory data into a form appropriate for use by the motor system. For example, in visually guided reaching the location of a target relative to the body is determined by a combination of the position of its image on the retina and the direction of gaze. What assures that the motor network responds to the appropriate combination of sensory inputs corresponding to target position in body- or arm-centered coordinates? To answer these questions, we model a sensory network coding target position and use it to drive a similarly modeled motor network. To determine the actual motor response we use decoding methods that have been developed and verified in experimental work. We derive a general set of conditions on the sensory-to-motor synaptic connections that assure a properly aligned and transformed response. The accuracy of the response for different numbers of coding cells is computed. We show that development of the synaptic weights needed to generate the correct motor response can occur spontaneously through the observation of random movements and correlation-based synaptic modification. No error signal or external teaching is needed during this process. We also discuss nonlinear coordinate transformations and the presence of both shifting and nonshifting receptive fields in sensory/motor systems.

[Key words: sensory-motor integration, neural coding, visually guided reaching, coordinate transformations, population decoding, correlation-based learning]

Reaching for an object, turning toward a visual or auditory cue, and a host of other sensory-guided motor tasks require the transfer of information from sensory to motor systems. Specifically, in these tasks sensory information about target location must be transferred to motor networks that generate movement toward the target. The coding of both target location in sensory net-

works (Knudsen and Konishi, 1978; Andersen, 1989; Takahashi, 1989; Stein, 1992) and movement direction in motor networks (Georgopoulos et al., 1986, 1988, 1993; van Gisbergen et al., 1987; Lee et al., 1988) involves the collective activity of large numbers of neurons individually tuned to respond over a broad range. Given this population coding strategy, how is the appropriate information transferred accurately from sensory to motor systems?

To produce accurately guided movement, a sensory network coding target location must, through synaptic connections, evoke a pattern of firing in a motor network that generates movement to the target. By sensory network we mean an array of neurons coding sensory information about target location at some late stage in a sensory pathway. Likewise, the motor network we discuss is a similar array coding the direction of movement at an early stage in a motor pathway. Three basic questions arise when we consider the flow of information between such neuronal arrays. The first is the alignment problem: what set of synaptic connections assure that activity in a sensory network coding target position evokes a response in the motor system that generates movement precisely toward the target? This question is often made more complex by a second issue, the transformation problem. Coordinate transformations may be required to put the information coming from sensory input into a form that is useful for generating motor output. Consider, for example, visually guided reaching. In such a task, the position of the image of the target on the retina does not indicate where the target is located relative to the head or body. Instead, a correction for the direction of gaze must be made and the motor system must respond to an invariant combination of retinal target position and gaze direction. Further transformations placing the target location in body- and arm-centered coordinates are required as well (Caminiti et al., 1990, 1991; Flanders et al., 1992; Kalaska and Crammond, 1992; Soechting and Flanders, 1992). What set of synaptic connections assures that the correct combination of sensory data (retinal position, gaze direction, head and arm position) is transferred to the motor network to evoke a pattern of motor activity that generates the appropriate motor response? Finally, a third basic question concerns the development and maintenance of these synaptic connections: what synaptic modification mechanisms assure that the correct synaptic connections arise during development and training and are maintained during ongoing activity?

The Model and Approach

It might appear that addressing these questions would require the construction of a full neural network model of sensory-guided reaching tasks (Bullock and Grossberg, 1988; Kawato et al., 1988; Kuperstein, 1988a,b; Jeannerod, 1990; Gaudio and Grossberg, 1991; Burnod et al., 1992; Kettner et al., 1993; Lu-

Received Feb. 2, 1995; revised Apr. 24, 1995; accepted May 26, 1995.

This work was supported by NSF-DMS9208206, the W. M. Keck Foundation, the McDonnell-Pew Center for Cognitive Neuroscience at Oxford (L.F.A.), and the Conacyt-Fulbright-IIIE program (E.S.). We thank A. Georgopoulos and C. Gross for informative and encouraging discussions.

Correspondence should be addressed to Emilio Salinas at the above address.
Copyright © 1995 Society for Neuroscience 0270-6474/95/156461-14\$05.00/0

kashin and Georgopoulos, 1993). However, by making use of what is known about the representation of target location and movement direction by the ensemble activity of arrays of neurons, we can both simplify the model building process and increase the accuracy of the model. We do not attempt to model how firing activity arises in the neurons coding sensory information, but instead use mathematical fits of average firing rate tuning curves. The firing responses within the sensory array in our model are generated from these tuning curves combined with randomly generated fluctuations. These allow us to include the high degree of variability seen in the firing rates of cortical neurons (Softky and Koch, 1992, 1994; Shadlen and Newsome, 1994).

Firing in the simulated sensory array drives a motor array through a set of synaptic connections. To determine the motor response generated by the resulting pattern of firing in the motor array, we use population decoding methods developed precisely for this purpose in experimental work. Population decoding has been used to relate the direction of arm movements of monkeys performing reaching tasks to firing activity in the motor cortex (Georgopoulos et al., 1986, 1988, 1993), premotor cortex (Caminiti et al., 1991), and cerebellum (Fortier et al., 1989). We use this and similar techniques (for a review, see Salinas and Abbott, 1994) to determine the movement that would be generated by a particular pattern of activity in the motor array of our model. This enables us to compute, for a given set of synaptic connections, the accuracy of the resulting movement and thus to determine what synaptic connections produce correctly transformed and aligned motor responses.

Important insights about how coordinate transformations may be performed by neural circuits have been provided by Andersen and collaborators (Zipser and Andersen, 1988; Andersen et al., 1993). Area 7a parietal neurons simultaneously represent both gaze direction and retinal position in a multiplicative manner (Andersen and Mountcastle, 1983; Andersen et al., 1985). Individual neurons in area 7a are tuned to the retinal position of a visual stimulus. When the gaze angle shifts, each tuning curve retains its retinal position selectivity but its height is modulated by a gaze-angle-dependent gain field. This allows both retinal position and gaze direction to be simultaneously encoded by the same sensory array. It has been suggested that such a representation is sufficient to allow downstream networks to extract the information needed to guide movement (Andersen, 1989; Stein, 1992; Andersen et al., 1993). In this view, a gaze-angle-invariant neuronal representation of target position does not have to appear anywhere within the sensory system. Zipser and Andersen (1988) used back-propagation methods to construct a network model with an intermediate representation like that in parietal area 7a and with an output providing a gaze-angle-invariant target position. However, as Zipser and Andersen (1988) themselves note, neural representations are likely to be distributed throughout the nervous system. Therefore, the crucial issue is not whether a mathematical neural network can generate an invariant output but rather whether the mixed representation of retinal position and direction of gaze can evoke distributed activity within a motor network that generates the correct motor response. Our model and approach allow us to examine this issue and they provide a mechanistic understanding of how distributed sensory and motor representations interact as well as a quantitative determination of movement accuracy.

We consider the transfer of information between only two networks, clearly a drastic simplification. Nevertheless, the two

network model allows us to examine both the transfer of target information and coordinate transformations of this information. Our approach and results can be applied to any neural system where a sensory array guides the generation of movement by a motor network. Examples include saccades to visual (van Gisbergen et al., 1987; Lee et al., 1988) or auditory (Knudsen and Konishi, 1978; Takahashi, 1989) targets or simpler actions like the bending reflex of the leech (Kristan, 1982; Lockery and Kristan, 1990a,b). Our analysis can also be extended to purely sensory systems (see Discussion). However, we will mainly focus on visually guided reaching tasks, in particular, one- and two-dimensional reaching tasks at fixed distance, for example, touching a spot of light on a screen. The coding of variable distances can be included using a similar approach (Pouget and Sejnowski, 1994). It has been suggested (Andersen, 1989; Glickstein, 1990; Stein, 1992; Andersen et al., 1993) that a representation of visual target location in area 7a is combined with arm position information in area 7b to generate a representation of intended movement direction in premotor and then motor cortex (Kalaska and Crammond, 1992). Along the way, a series of coordinate transformations are made from retinal to head-, body-, and arm-centered coordinates (Caminiti et al., 1990, 1991; Flanders et al., 1992; Kalaska and Crammond, 1992; Soechting and Flanders, 1992). We will focus on the transformation from target location in retinal coordinates to the representation of movement direction in head-centered coordinates. The same mechanisms that we uncover in this example can be used to generate transformations to body- and arm-centered coordinates that incorporate the effects of head (Biguer et al., 1988; Roll et al., 1991; Brochier et al., 1995) and arm position (Caminiti et al., 1990, 1991).

We call the location that is the final target of an arm movement the goal location. Correct alignment of the sensory and motor representations occurs when the target location and movement goal location are identical. For our model, we will derive a condition on the weights of the synapses coupling the sensory array to the motor network that guarantees accurate alignment of the motor response to a correctly extracted combination of sensory signals. Synaptic weights satisfying this condition arise spontaneously in the model from an unsupervised correlation-based learning rule when, during development, the motor system provides the stimulus for the sensory system. This form of learning has been proposed in childhood development (van der Meer et al., 1995) and has been used in other models of reaching tasks (Kuperstein, 1988a,b; Gaudio and Grossberg, 1991). In our model, development of the correct synaptic connections by this mechanism occurs spontaneously through observation of random movements and does not require any externally or internally generated error signals.

By using prefit tuning curves and population decoding techniques we avoid considering how the sensory array finds and identifies the target or how the transformations from a coded direction to limb movement kinematics and trajectories (Abbott and Blum, 1995) and force dynamics (Alexander and Crutcher, 1990a,b; Crutcher and Alexander, 1990; Mussa-Ivaldi and Giszter, 1992; Redish and Touretzky, 1994) are made. These are interesting and important issues but it is a strength of our approach that we can address other relevant questions without having to deal with these problems as well.

The first example we study with our model and methods is the transfer of target location from a sensory to a motor array without including any coordinate transformations. In a visually guided task, this corresponds to the transfer of information from

a sensory network coding retinal position of a target at fixed gaze angle, to a motor network coding movement direction. In this case no distinction needs to be made between target location and retinal position. We then include the effects of a variable gaze angle to study how the correct coordinate transformation is made. Finally, we return to the case of fixed gaze direction to study the simultaneous encoding of many variables and to demonstrate a nonlinear coordinate transformation between the sensory and motor arrays.

Methods

We apply a combination of analytic and computer simulation techniques to study the problem of control of motor actions by a sensory array of neurons coding target position. The analytic approach requires some simplifying assumptions and approximations but produces very general results. These are then tested and extended by computer simulations that do not rely on these simplifications.

List of symbols

x = the location of the target coded in the sensory array. When we consider variable gaze directions, this is the target location relative to the fixation point or, equivalently, the retinal position of the target. For one-dimensional examples, x is a single number while in two-dimensional cases it is a vector with Cartesian components (x_1, x_2) .

y = the location of the fixation point characterizing the direction of gaze.

z = the goal location of the movement generated by the motor array. In some of our examples this is instead an angle characterizing the movement direction.

R_i^s = the firing rate of neuron i in the sensory array.

R_i^m = the firing rate of neuron i in the motor array.

η_i^s = a random noise term included in the response of neuron i in the sensory array.

η_i^m = a random noise term included in the response of neuron i in the motor array.

a_i = the target location evoking the maximum average firing response in neuron i of the sensory array (at fixed gaze direction when this is included). We call this the preferred target location or retinal position. Like the target location it is either a single number or a two-dimensional vector.

b_i = a particular gaze fixation point characterizing the response of neuron i in the sensory array as a function of gaze direction for a fixed retinal target location. We call this the preferred location of gaze although, for the linear gain fields we use, the average response of neuron i actually reaches its maximal plateau at $y = b_i$ (see below).

c_i = the movement goal location associated with the maximum average firing rate of neuron i in the motor array. We call this the preferred movement goal location for neuron i .

$f(|a_i - x|)$ = the average firing rate of neuron i in the sensory array when the target location is x and the gaze direction is fixed. The notation $|a_i - x|$ means that the average firing rate depends on the distance between the actual target location and the preferred target location for this neuron.

$f(|a_i - x|, |b_i - y|)$ = the average firing rate of neuron i in the sensory array when the target location in retinal coordinates or retinal position is x and the gaze direction is y .

$g(|c_i - z|)$ = the average firing rate of neuron i in the motor array when the movement goal location is z .

W_{ij} = the strength of the synaptic connection from neuron j in the sensory array to neuron i in the motor array.

We use the notation $[R]_+ = R$ if $R \geq 0$ and zero otherwise for any quantity R . This is used to eliminate negative firing rates. In addition, $[R]_M^+ = M$ if $R \geq M$ and $[R]_+$ otherwise. This is used to set a maximum firing rate M .

Firing rate tuning curves and variability

We reproduce the firing activity of neurons in the sensory array responding to a particular target location by modeling the average firing rates of the coding neurons and their variability about this mean. The average firing rates are characterized by tuning curves that are functions of target location. When we consider a sensory array coding both retinal position and gaze direction, the average firing rate tuning curves depend

on both of these quantities. For simplicity, tuning curves for different neurons in a given array are assumed to have the same shape but different preferred locations. The average firing rate tuning curves we use can have almost any form but they must have one important property: the average firing rate for a coding neuron depends only on the distance between the location being coded and the preferred location for that neuron. Mathematically, for the sensory array at fixed direction of gaze with target location x , the firing rate of a neuron with preferred target location a_i will be

$$R_i^s = f(|a_i - x|) + \eta_i^s. \quad (1)$$

The function f is the average firing rate and η_i^s is a random variable that introduces firing rate fluctuations. Random fluctuations for the different neurons are generated independently from a Gaussian probability distribution with zero mean and standard deviation equal to the average firing rate. Thus, the firing rate variability is characterized by a dimensionless coefficient of variation equal to one (Softky and Koch, 1992, 1994). We have chosen Gaussian rather than Poisson variability so that the fluctuations can be characterized by a dimensionless parameter. Modeled in this way, firing rate fluctuations are uncorrelated between neurons. The effect of correlated noise is considered in the Discussion.

When a variable gaze direction y and preferred gaze location b_i are included, the firing rate is expressed as

$$R_i^s = f(|a_i - x|, |b_i - y|) + \eta_i^s. \quad (2)$$

In this case, x is the retinal position of the target and a_i is the preferred retinal position for neuron i .

The assumptions that we have made about the average firing rate tuning curves may seem more restrictive than they really are. The tuning curves of individual neurons do not have to be identical, symmetric, or even peaked for our results to be valid. This is because our results still apply if we replace the tuning curves of individual neurons with average tuning curves obtained from summing a number of neurons with similarly placed receptive fields. Thus, the functions f (and g defined below for the motor array) can be thought of as average tuning curves obtained from neurons with similar preferred locations. In this average it is possible to combine tuning curves that individually differ, are asymmetric, and have no preferred location and still obtain average tuning curves with all the properties we assume. In the case of the linear gain curves characterizing gaze direction effects, we will make use of this trick. We should stress that our proofs can be extended to the case of nonsymmetric and nonidentical tuning curves if assumptions about uniformity across the population of coding neurons are made. However, to keep the derivations simple we do not include these extensions.

For the analytic work we use average firing rates. The effects of fluctuations are considered in the computer simulations. This is because the fluctuations we include are too large for the usual linear approximations used to calculate their effects analytically to be valid. The analytic results are extremely general, they apply for any functions f and g (see below) characterizing the average neuronal responses. For computer simulations we must, of course, choose definite functions for the neuronal tuning curves. For the sensory array we use Gaussian tuning curves, which fit the retinal position data from parietal cortex (Andersen et al., 1985; Zipser and Andersen, 1988). At fixed gaze direction this means

$$f(|a_i - x|) = R_{\max} \exp\left(-\frac{|a_i - x|^2}{2\sigma^2}\right). \quad (3)$$

If the width of these curves is defined at $e^{-1/2}$ of their peak value R_{\max} , this tuning curve has a width of 2σ . When variable gaze direction is included we express the tuning curves as the product of a Gaussian function of the retinal position of the target x and a gain field that depends on the direction of gaze y . The gaze direction effects found in parietal area 7a were described as products like this with approximately linear gain modulation (Andersen et al., 1985). A fully linear gain modulation is awkward for us because it is unbounded, so instead we use a clipped linear function that cannot increase beyond a certain maximum value M . Thus, as a function of both retinal position and gaze direction,

$$f(|a_i - x|, |b_i - y|) = \frac{R_{\max}}{M} \exp\left(-\frac{|a_i - x|^2}{2\sigma^2}\right) [\pm(b_i - y) + M]_+^M, \quad (4)$$

where the notation $[]_+^M$ defined in the List of symbols means the gain modulation is linear but bounded between 0 and M so the averaging

firing rate is between zero and R_{\max} . The plus/minus sign allows for gain fields that either decrease or increase as a function of gaze direction. The preferred gaze location b_j marks the points where the maximum value of the gain field is reached. Note that these tuning curves are not functions of the absolute value of $b_j - y$. However, if we average the tuning curves from two neurons with the same preferred gaze location b but one with a plus sign and one with a minus sign in Equation 4, the combination is a function of $|b_j - y|$ as in Equation 2. As mentioned previously, use of average tuning curves allows us to apply the theorems proved in the Appendices without adding additional complications, although it is not essential.

In some of our studies we examine the coding of two or more quantities by a single sensory array. In these studies we express the average firing rate as a product of Gaussian tuning curves for each of the coded quantities. For example, when we consider the coding of target location in two dimensions we use tuning curves like Equation 3, where x and a_i are two-dimensional vectors and the absolute value represents the vector length.

To decode the population response of the motor array, we need to relate the average firing rates of the neurons to the direction of movement. If the movement is directed toward the location z , we assume that a neuron in the motor array with preferred movement goal location c_i fires at an average rate $g(|c_i - z|)$ that depends on the distance between these two locations. For the motor system we use tuning curves that are either Gaussian,

$$g(|c_i - z|) = R_{\max} \exp\left(-\frac{|c_i - z|^2}{2\sigma^2}\right), \quad (5)$$

or cosine,

$$g(|c_i - z|) = R_{\max} \left[\cos\left(\frac{\pi|c_i - z|}{2\sigma}\right) \right]_+, \quad (6)$$

as seen in the motor cortex of the monkey (Schwartz et al., 1988). The notation $[]_+$ defined in the List of symbols prevents negative firing rates. In addition to eliminating the negative portion of the curve, we also use only one cycle of the cosine function (see Fig. 1). The width of the cosine tuning curves is measured between the zero crossings and is 2σ . Note that the average firing rate tuning curves for the motor array are not used to generate firing rates (except in our initial study of decoding accuracy) but to decode them.

Synaptic weights and modification rules

In our studies, the activity in the sensory array that encodes target position is used to drive the motor array. This is done through synaptic coupling between the sensory and motor networks. We characterize the strength of the synaptic connection from neuron j of the sensory array to neuron i of the motor array by the weight W_{ij} . The firing rate of neuron i in the motor array is given by

$$R_i^m = \left[\sum_j W_{ij} R_j^s \right]_+ + \eta_i^m, \quad (7)$$

where the $[]_+$ notation again eliminates negative firing rates. The term η_i^m introduces random Gaussian fluctuations with standard deviation equal to the mean firing rate just as in the case of the sensory array.

We will show that a simple training procedure allows the synaptic weights in Equation 7 needed to produce accurate movement toward a sensory-coded target to develop spontaneously. Initially we considered two possible candidate rules for computing how activity modifies synaptic strength. In both cases the synaptic weights arising from a development or training period depend on the average correlation between pre- and postsynaptic firing rates during that period. We denote averages over the development period by brackets $\langle \rangle$. For the Hebb rule (Hebb, 1949), the strength of the synapse from neuron j in the sensory array to neuron i in the motor array following the development/training period is

$$W_{ij} = \langle R_i^m R_j^s \rangle. \quad (8)$$

(We could include a proportionality constant in this equation but, for our purposes, this makes no difference. Our results depend on the pattern of synaptic weights, not on their absolute magnitude.) Another alternative is the covariance rule (Sejnowski, 1977), for which

$$W_{ij} = \langle R_i^m R_j^s \rangle - \langle R_i^m \rangle \langle R_j^s \rangle. \quad (9)$$

In our case, the last term on the right side is a constant, so we use a rule that encompasses both cases,

$$W_{ij} = \langle R_i^m R_j^s \rangle - k, \quad (10)$$

and we generalize it by allowing k to be an arbitrary constant. In the results shown, k was set to optimize the performance of the system.

The training procedure we use consists of generating random movements that act as target locations for the sensory array. In the case when the gaze direction is fixed, this means that the target location x is equal to the goal location z because the sensory target is the moving arm itself. During the development/training period random motor activity positions the arm at various different goal locations. We assume that if the arm is observed during this process, the movement goal location will be represented in the sensory array as a target location. Random movement corresponds to varying z uniformly over its entire range and observation of these goal locations causes the target location x to equal z . During these random movements, the firing rates in the motor array will, on average, be given by $g(|c_i - z|)$ and those of the sensory array by $f(|a_j - x|)$ with $x = z$. If the synaptic weights develop according to the rule (Eq. 10), the resulting pattern of synaptic weights that arises is given by the integral

$$W_{ij} = \int dz g(|c_i - z|) f(|a_j - z|) - k. \quad (11)$$

Average firing rates can be used here because the added noise terms are independent and average to zero. These weights are used when we study the case of fixed gaze direction.

When the gaze direction is allowed to vary, the training procedure is identical except that the gaze direction y is randomly varied along with the movement goal location z . Again we assume that the goal location of the randomly moving limb provides the target location for the sensory system. The target location in head-centered coordinates corresponding to retinal position x and gaze direction y is $x + y$. The fact that the movement goal location is providing the sensory target during training therefore imposes the condition $x + y = z$. With this constraint and the arm movement goal location and gaze direction randomly varying, we have integrals over the full ranges of y and z with $x = z - y$ so that

$$W_{ij} = \int dy dz g(|c_i - z|) f(|a_j - z + y|, |b_j - y|) - k. \quad (12)$$

These weights are used in our studies involving a variable direction of gaze.

Decoding methods

To interpret the firing responses produced in the motor array of our model, we must determine what movement goal location they represent. We do this in two different ways. The first is a variant of the approach developed in studies of monkey motor cortex (Georgopoulos et al., 1986, 1988). This is the vector method, in which the decoded goal location of the movement is a firing-rate weighted average of the preferred locations of the individual neurons:

$$z = \frac{\sum_i R_i^m c_i}{\sum_i R_i^m} \quad (13)$$

We use this method of decoding in all of our analytic work and in some of the simulations. However, we use a more accurate decoding technique in many of the simulations. This is the maximum overlap technique. Here the decoded value of z is the movement goal location that maximizes the overlap between the observed firing rates and the average firing rates so that

$$\sum_i R_i^m g(|c_i - z|) = \text{maximum}. \quad (14)$$

In the cases where we use this method it is equivalent to a least squares fit of the observed rates to the average rates (Salinas and Abbott, 1994). All of the analytic results presented here using the vector decoding method can be reproduced using the maximum overlap technique. However, the proofs and derivations are more complex for this latter case so we only present the simpler vector method results. Likewise, there are no significant differences between the simulation results we obtain using the two methods. We present results obtained from the most appropriate method for a particular example.

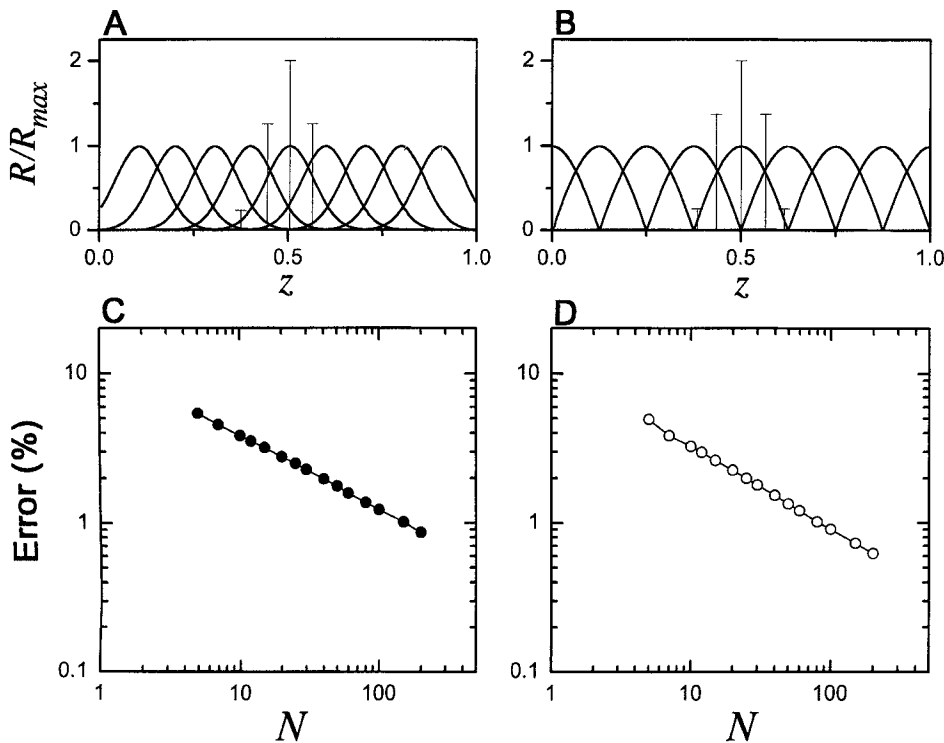


Figure 1. Decoding accuracy using the maximum overlap method. The coded variable z lies between zero and one. For this example, no jitter was included in the positioning of preferred locations for the coding neurons so that *A* and *B* were clearer. *A*, An array of Gaussian average firing rate tuning curves with the standard deviation of firing rate fluctuations indicated. The width of the Gaussian curves is $1/8$. *B*, A similar array of cosine tuning curves with width equal to $1/4$. *C*, The decoding accuracy for the Gaussian array as a function of the number of neurons N in the array. Accuracy is the rms decoding error expressed as a percentage of the full range of z . *D*, Same as *C* but for the cosine array.

Analytic techniques

To perform the analytic calculations we use an extremely useful approximation, the replacement of sums over neurons by integrals over their preferred locations. We assume that the coverage of locations by the coding neurons is dense and uniform. This allows us to approximate sums with integrals. We have checked the accuracy of this approximation by computer simulation and it is remarkably good even for modest numbers of neurons. We ignore the limits in the resulting integrals by allowing the coded quantities to vary over an infinite range. In the simulations, these ranges are obviously finite but we avoid working too close to the limits.

Computer simulations

In the computer simulations, sensory and motor arrays of a given size were constructed by assigning preferred locations to the different neurons. Preferred locations were determined by laying out a uniform grid over the range of coded values and then introducing a random variation about these grid locations. We call these shifts "jitter" and chose them randomly between plus and minus one-quarter the width of the corresponding tuning curves. We have obtained similar results by choosing completely random preferred locations, although this increases the size of the errors somewhat. In some cases we used zero jitter for clarity.

We began a typical simulation by choosing a target location (or retinal position and gaze direction) randomly. Firing rates in the sensory array corresponding to this target location (and gaze direction) were generated as the sum of a mean firing rate and a randomly generated fluctuation. When this sum was negative the fluctuation was discarded and a new one was chosen randomly.

The sensory responses were used to drive a motor array using Equation 7. Again, negative rates were avoided by selecting another random fluctuation. The activity generated in the motor array was decoded as outlined above. Maximization of the overlap was done using Brent's method (Press et al., 1992). Movement error was computed by comparing the decoded goal location of the movement with the target location used to drive the sensory array. A range of target locations or, in later examples, retinal positions and gaze directions was sampled randomly. A large number of runs were performed to accumulate good statistics on the size of the average root-mean-square (rms) error in the motor response.

Results

Decoding accuracy

Before analyzing coupled sensory and motor networks, we will show that the decoding procedure we use to determine the movement direction resulting from a particular pattern of activity in the motor array is accurate. To test decoding accuracy, we considered arrays of neurons with either Gaussian (Eq. 5) or cosine (Eq. 6) tuning curves. We randomly chose a "true" movement direction and generated a set of firing rates using the tuning curves of the motor array to set the mean values and including random fluctuations. We then decoded those responses using the maximum overlap method. Accuracy was tested by comparing the decoded movement direction with the original direction used to generate the rates. The results are shown in Figure 1, where *A* and *B* illustrate the arrays of tuning curves we used as well as the standard deviation of the firing-rate fluctuations included in the model, and *C* and *D* show the decoding errors for arrays with different numbers of neurons. The decoding error is the rms average (expressed as a percentage) of the difference between the decoded and actual goal locations divided by the full range of possible goal locations, which is equal to one in this example. Figure 1 indicates a high degree of decoding accuracy. For more than about 20 coding neurons, the decreasing error is proportional to one over the square root of the number of neurons, indicating that the method has reached the point where noise effects are the primary limitation on the decoding accuracy.

The accuracy of the decoding is quite remarkable considering the large variability in the firing rates of the model neurons. Figure 2 shows this variability for the entire population of neurons on one typical run. For a fixed value of z , the firing rate of each cell in a 200 neuron array is plotted as a function of its preferred location. The dashed curve shows the mean firing rates for these neurons. Although the fluctuations about the mean are

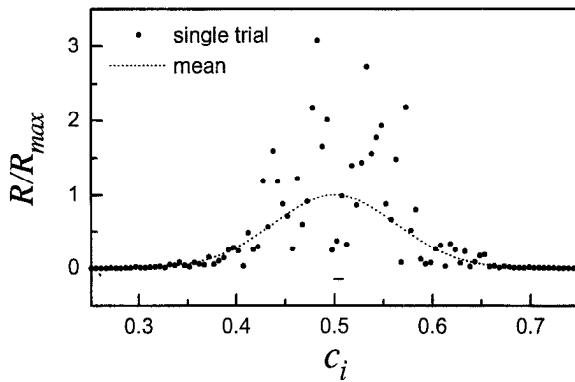


Figure 2. Decoding accuracy for a single trial using the maximum overlap method. This shows a single trial using an array like that of Figure 1A. The firing rate of each neuron in a 200 cell array is plotted at a position on the horizontal axis corresponding to its preferred location. The dashed curve shows the average firing rate. Despite the large fluctuations about this average the decoding error, indicated by the small bar near $c_i = 0.5$, is small. The actual encoded value was 0.5 and the decoded value was 0.507, an error of 0.7%.

large, the error in the decoded value, as shown in the figure, is remarkably small. For this to occur, it is most important that silent neurons or those firing at low rates show little variability while large variability in rapidly firing neurons is less harmful to decoding accuracy. This is, of course, exactly a property of Gaussian fluctuations with standard deviation proportional to the average firing rate or of Poisson firing rate fluctuations that have the variance proportional to the average rate. We have obtained similar results using other types of tuning curves, such as sigmoids or more complex, multimodal shapes. We have also found that accurate decoding results can be obtained even if the tuning curves used in the decoding procedure are not exactly those used to generate the firing rates. The tuning curves used for decoding must have preferred locations fairly close to the actual tuning curves but the shapes or widths do not have to match particularly precisely.

Alignment of sensory and motor networks

Having checked decoding accuracy, we now use a sensory network to drive the motor network that we decode. In this case, the firing rates of the neurons of the motor array were not computed from their average tuning curves as they were for Figure 1. Rather, the motor array was driven by the sensory array according to the rule (Eq. 7). In this section we consider the case of fixed gaze direction. The rates in the sensory array for a randomly chosen target location were generated from the average firing rate tuning curves of the sensory array along with random fluctuations. The direction of movement was determined from the firing rates of the motor array as given by Equation 7 through decoding. The goal location of the decoded movement was then compared with the target location used originally to generate firing in the sensory array.

What properties of the synaptic weights assure that firing in the sensory array corresponding to a particular target location will evoke activity in the motor array that produces movement toward the target? In Appendix A we derive a simple condition that assures that this alignment will occur for any tuning curves that are functions of the distance between the actual and preferred target or movement locations. The condition is simply that the strength of the synaptic connection from neuron j to neuron

i depends only on the magnitude of the difference between their preferred locations. Mathematically,

$$W_{ij} = W(|c_i - a_j|). \quad (15)$$

Note that a wide variety of functions W can be used as long as they depend in this way on the two preferred locations. Remarkably, this condition is all that is required to assure that the two networks align. Equally remarkable is the fact that synaptic connections with the necessary property can be produced by a simple, unsupervised training process and correlation-based synaptic learning rule.

The procedure that produces an appropriate set of synaptic couplings contains one essential element, the goal location of the motor array must provide the target coded by the sensory array during the time that the synaptic weights develop. In the case of visually guided arm movements, the learning procedure would consist of randomly waving the arm around and watching it so that its position is coded in the sensory array. In Appendix B we show that the synaptic weights (Eq. 11) that arise from this process have exactly the property (Eq. 15) needed to align the two networks.

The results of computer simulations based on this idea are shown in Figure 3. To produce these results we coupled a sensory and a motor array with synaptic weights generated by correlation-based synaptic modification during observation of random movements as given by Equation 11. In Figure 3, the accuracy of the resulting movement direction is plotted for different array sizes. Figure 3A shows a one-dimensional example much like that shown in Figure 1C, but here the motor array is driven by a sensory array. The accuracy is comparable to the decoding accuracy shown in Figure 1, indicating a nearly perfect transfer of target location from the sensory to the motor array. As indicated in Figure 3A, this type of information transfer does not require that the sensory and motor arrays have the same types of firing-rate tuning curves. We have verified that accurate alignment occurs when the sensory array has Gaussian tuning curves and the motor array has either Gaussian tuning curves with different widths from those of the sensory array, as shown in Figure 3A, or cosine tuning curves (see, e.g., Fig. 7).

Clearly synaptic weights generated by watching random movements and performing correlation-based synaptic weight modification produce extremely good alignment between the networks. Figure 3B shows an application of this idea to a situation closer to actual two-dimensional reaching experiments. Here a two-dimensional target location vector evoked a response in a sensory array that drove a motor array through a set of synaptic weights arising from the same training process. However, instead of coding target and movement goal location we assumed that the sensory and motor arrays coded angles representing the direction to the target and the direction of the resulting movement. The tuning curves used were cosines of the appropriate direction angle. The error plotted is the angular difference (in degrees) between the target direction evoking a response in the sensory array and the movement direction inferred by decoding the activity of the motor array generated by that response. For networks of 100 neurons, the error is only a few degrees and it falls as $N^{-1/2}$ with increasing numbers of network neurons.

Coordinate transformations between sensory and motor networks

As noted in the introduction, sensory information often has to be transformed before it can correctly guide motor actions. In

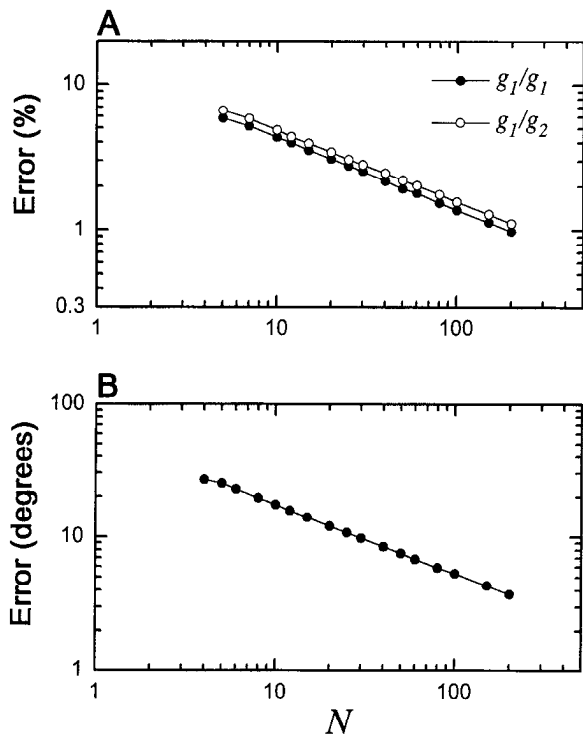


Figure 3. Accuracy of information transfer from a sensory to a motor array. A sensory array coding target location at fixed gaze direction drove a motor array which was then decoded using the maximum overlap method. Synapses were those that developed during random "watched" movements according to the correlation-based learning rule (Eq. 10) with $k = 0.055$. Jitter in the positioning of individual tuning curves was included for both arrays. Both arrays had the same number of neurons, N , which was varied. *A*, g_1 and g_2 correspond to two Gaussians of widths $1/8$ and $1/(8\sqrt{2})$, respectively. *Filled circles* show the transfer from a sensory array to a motor array using Gaussians g_1 of equal width. *Open circles* show the transfer from a sensory array with Gaussian tuning curves g_1 to a motor array with Gaussians g_2 . *B*, Information transfer between a sensory array coding direction to a target and a motor array coding the direction of the evoked movement. Both arrays use cosine tuning curves with widths equal to π . Jitter was set equal to ± 0.16 radians. The value $k = 0.1$ was used in the synaptic learning rule. The vector method of decoding was used and the error reported is the rms angular difference between the target and movement directions.

the case of a sensory array coding both retinal position x and gaze direction y , the motor array must respond to the appropriate combination of these two quantities. For the definition of retinal position that we use, the movement required to reach the target must be toward the goal location $z = x + y$. As in the case of fixed gaze direction, the motor response must align properly but, in addition, it must align to this particular combination of the two directions simultaneously coded in the sensory array.

In Appendix C we derive a condition on the synaptic weights that guarantees that the activity of a sensory array coding both retinal position and gaze direction will evoke activity in a motor array that generates the correct movement to the coded target. This condition is again remarkably simple. We assume that the retinal position and gaze direction are simultaneously encoded in a single array but allow any form of simultaneous encoding. In particular, we do not require multiplicative coding or linear (planar) gain fields, we only require that Equation 2 is satisfied with any function f . In this case, neurons in the sensory array are characterized by two preferred locations, one for retinal po-

sition and one for gaze direction. The correct motor response is assured if the synaptic weights coupling the sensory to the motor array depend only on the magnitude of the difference between the preferred direction of the motor neuron and the sum of the preferred retinal position and gaze direction for the sensory neuron. In other words, the motor response will be $z = x + y$ provided that (Appendix C)

$$W_{ij} = W(|c_i - a_j - b_j|). \quad (16)$$

Again, the function W is not unique.

Synaptic weights satisfying Equation 16 can be generated by a correlation-based synaptic modification rule through a procedure similar to the one used in the case of fixed gaze direction. The only addition is that the gaze direction must be varied randomly while the motor array randomly stimulates the sensory array during development of the synaptic weights. In Appendix D, we show that the synaptic weights (Eq. 12) produced by this procedure satisfy Equation 16. In simple terms, this means that if the gaze direction is varied while the arm is waved around randomly acting as a visual target, the correct gaze-angle-invariant alignment will take place automatically.

Figure 4 shows the results of a simulation based on synaptic weights generated by this mechanism. Here the sensory cells were tuned to the retinal position of the target, with response amplitudes modulated by a gain field that was linear up to a saturation value, beyond which it remained constant. The sensory responses had the form given by Equation 4, with equal numbers of increasing and decreasing gain fields included in the array. After determining the synaptic weights according to the correlation-based rule (Eq. 12), the sensory responses were used to drive the motor array that coded the goal location of the movement. The error between the decoded goal location of the movement and the correct target location $z = x + y$ was then obtained. This is plotted for different values of z in Figure 4A. Each point is the average of several combinations of retinal target positions and gaze directions corresponding to the same target location in head-centered coordinates. The error is a small fraction (expressed as a percentage) of the range of values that the movement goal location can take. Figure 4C shows that the error decreases as one over the square root of the number of cells in the network.

It is interesting to note that individual neurons of the motor network in this example have shifting receptive fields. This occurs even though the sensory neurons driving them have fixed receptive fields that are gain-modulated by gaze direction but not shifted. In Figure 5A we plot the average response of a neuron in the sensory array as a function of retinal target position for several gaze directions. The tuning curves are modulated but they do not shift. Figure 5B shows average responses of a neuron in the motor array plotted as a function of the retinal target position evoking them for different gaze directions. The magnitude of the shift precisely matches the change in gaze direction. These tuning curves have retinal receptive fields that shift as a function of gaze angle but that are fixed as a function of target location in head-centered coordinates.

A slight extension of the proofs in Appendices C and D allows us to generalize these results. When a sensory array coding two locations x and y drives a motor network, the movement response will be $z = \alpha x + \beta y$ for any values of α and β provided that

$$W_{ij} = W(|c_i - \alpha a_j - \beta b_j|). \quad (17)$$




Figure 4. Information transfer from a network coding both retinal target position x and gaze direction y to a motor network with decoded goal location z . The coded quantities x and y ranged from -10 to 10 . However, to avoid edge effects the full range was not used. For the results shown, variables were restricted to $-6.5 < x, y < 6.5$ and $-4 < z < 4$. The value $k = 15$ was used in the synaptic learning rule. All tuning curves were Gaussian with width equal to 2. Gain fields were clipped-linear saturating at $M = 3$. In *A* and *B* no jitter was present but it was used for *C*. In *A* and *B*, the sensory array had 676 neurons, the motor array 169. *A*, The percentage error in the goal location of the coded movement relative to the true target location, which is $x + y$. This is plotted over a range of goal location values averaged over eight x and y combinations that give a particular fixed value of $x + y$. The solid curve shows the error with noise, and the dashed curve, the error without noise. Synapses were those arising from a correlation-based learning rule acting during the observation of random movements with random gaze directions. *B*, Same as *A* except that during the synaptic development the retinal target position was set equal to the goal location. As a result the motor response became aligned with x instead of with $x + y$ as in *A*. The error shown is the average difference between z and x . *C*, The average over z of the errors shown in *A* and *B* plotted for different network sizes. Filled circles show that case when z was correlated with $x + y$ and open circles correspond to z correlated with x . N is the number of neurons in the sensory array and the motor array has one-fourth this number of cells.

The case of retinal position and gaze direction corresponds to the particular case $\alpha = \beta = 1$. This means that any linear combination of x and y can be extracted from the sensory array by a downstream network. We show this in Figure 4. The same sensory responses as in Figure 4*A* were used to drive a different motor array that generated movement toward a goal z , which was simply the target location in retinal coordinates, x . This corresponds to the case $\alpha = 1, \beta = 0$ and amounts to generating the same type of movement as in Figure 3*A* but with the gaze direction modulation present in the sensory responses. The synaptic connections were generated by a rule similar to Equation 12, except that the constraint $x = z$ was imposed. In other words, during development the retinal position of the target was set equal to the goal location of the random movements. Then the movement goal location was decoded from the evoked activity in the motor array and compared to x . The error that decreases like one over the square root of the number of coding neurons is shown in Figure 4, *B* and *C*. Similar results were obtained with other combinations of α and β . Thus, a motor network receiving input from a sensory array simultaneously encoding two variables x and y can generate a movement to a goal location that can be any linear combination of x and y . A correlation-based rule for synaptic development along with observation of random movements will produce the appropriate synaptic weights. If a motor network generates movement that is correlated with a particular linear combination of x and y it will automatically become aligned with that combination of sensory inputs.

The results we have derived and shown also apply if the quantities x, y , and z are vectors and they do not depend on the gain fields being linear. We have confirmed this by carrying out the same kinds of simulations using Gaussian gain fields. The results are qualitatively the same: any linear combination $z = \alpha x + \beta y$ can be read out by the motor array; the error decreases as $N^{-1/2}$ for N neurons and the tuning curves of the motor array neurons, as functions of the sensory variable x , shift for different y values.

Representation of more than one sensory variable

The results of the last section raise the interesting subject of neuronal arrays that encode several quantities simultaneously. Neurons in cortical circuits may respond to a wide variety of stimuli, so many quantities might be encoded by their firing. In the last section we considered the case of simultaneous coding of retinal position and gaze direction. Neurons coding the location of auditory stimuli in the barn owl similarly display simultaneous coding of the two direction angles (Knudsen and Konishi, 1978; Takahashi, 1989), and numerous other examples exist (e.g., see Knudsen, 1982). To study simultaneous coding of multiple variables, we considered a sensory array of neurons that responded to a number of quantities with average firing rate tuning curves that were the product of Gaussian factors for each of the individual quantities. We then looked at the accuracy with which one of these quantities could be extracted by the motor array when a number were coded simultaneously. The result is shown in Figure 6. The error increases gradually with the number of variables encoded. This increase in the error does not

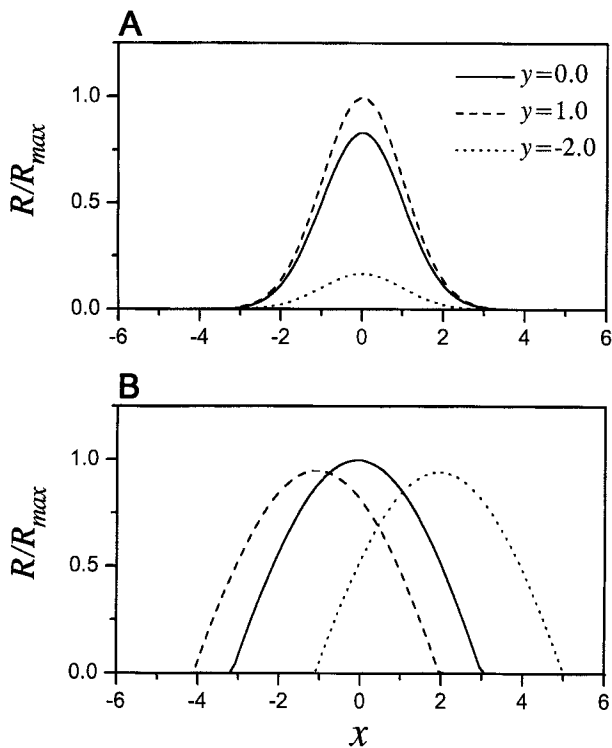


Figure 5. Nonshifting and shifting average firing rate tuning curves. **A**, The average firing rate of a neuron from a sensory array coding both retinal target position and gaze direction plotted as a function of retinal target position. Responses for three different values of gaze direction are plotted. The neuron has a retinal receptive field that does not shift but instead is modulated as a function of gaze direction. **B**, The average evoked response in a neuron of the motor array producing movement to a goal location $x + y$ plotted as a function of retinal target position for the same three gaze directions. The retinal receptive field now shifts and tracks the target location in head-centered coordinates. The tuning curves, networks and synapses used here are identical to those of Figure 4A.

necessarily degrade the performance of the system because, for each point, the error still decreases as one over the square root of the number of neurons. The rule used for the synaptic connections, Equation 11, was the same as for previous cases with one of the coded variables playing the role of target location. This variable was set equal to z during the development of the synaptic weights and all the other encoded variables were uncorrelated with z . Figure 6 shows that accurate alignment between sensory and motor responses can arise even when the sensory neurons code for a number of additional quantities that do not carry information relevant to the motor actions.

Nonlinear coordinate transformations

We have seen that a downstream network coupled with appropriate synaptic connections to an array of neurons coding two locations x and y can extract the sum $x + y$. This generates a representation of the location corresponding to x in a coordinate system shifted by y and provides a general mechanism for performing coordinate translations. However, other types of coordinate transformations may be desirable. Here we consider a two-dimensional example where target location is characterized by x_1 and x_2 , which are Cartesian coordinates. We represent the two-dimensional point (x_1, x_2) in a sensory array that simultaneously codes these two quantities with tuning curves that are Gaussian as in Equation 3. We ask whether the motor array can

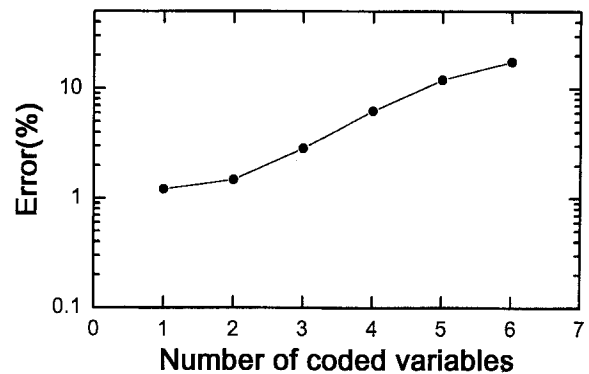


Figure 6. Accuracy as a function of the number of simultaneously coded quantities. Target location x was represented by an array of Gaussian tuning curves as in Figure 1A. A variable number of other quantities were represented by multiplying these tuning curves by randomly placed Gaussians of the other coded variables. The preferred location for x was distributed uniformly within its range. The sensory array had 5000 neurons and the motor array 200. The error between the motor response z and the target location x is plotted as a function of the number of quantities being simultaneously encoded along with x in the sensory array.

extract from this representation of target position the angular location of the target in polar coordinates. Thus, we consider the nonlinear transformation from Cartesian to polar coordinate systems rather than a simpler coordinate translation. This means that the motor response will be characterized by an angle z specifying the direction of movement. Correct alignment requires that this movement direction angle be equal to the target direction angle, $\theta = \arctan(x_2/x_1)$; that is, the network must effectively compute an arctan function. In this example, the tuning curves for the neurons of the motor array were cosines as in Equation 6. The synapses between these two arrays were constructed by the same correlation-based rule we used in the other examples. During the development process random movement in a particular direction was interpreted by the sensory array as a sequence of goal locations at various radial distances but with polar angle $\theta = z$. After synaptic development, the sensory array coding target location evoked activity in the motor array and we compared the decoded movement direction z with the polar angle θ corresponding to the target.

Figure 7 shows the results. The error depends on the length of the target position vector $r^2 = x_1^2 + x_2^2$. When r is small θ cannot be determined accurately and, not surprisingly, the error is large. As r increases the error decreases until we reach values of x_1 and x_2 outside the receptive fields of the sensory neurons where, naturally, it increases again. The inset shows that for a given value of r the error is the same, on average, for all directions. For the network of 500 neurons shown in Figure 7 the error is a few degrees. We have seen that this error also decreases as $N^{-1/2}$.

Figure 8 shows three response tuning curves of neurons in the motor array as functions of the polar coordinate angle of the target location θ for three different values of r . It is interesting to see that the polar coordinates of the target location are coded in this array much like retinal position and gaze direction are coded in parietal area 7a. Individual neurons have receptive fields that depend on θ and these responses are gain modulated by a factor that depends on r . This modulation is consistent with the error dependence on r . For constant variance, the decoding accuracy increases with the amplitude of the tuning curves (Ab-




Figure 7. Nonlinear coordinate transformation between the sensory and motor arrays. The sensory array responded to the Cartesian coordinates of the target location. The angular position of the target was correlated during training with the coded motor response z and then the error between the polar angle and the evoked value of z was determined. The main figure shows this error as a function of the length of the target location vector and the inset shows the error for a particular length, $r = 7$, over all movement direction angles. The *solid curve* in the inset is the error with noise, and the *dashed line*, the error without noise. The target location variables were in the range $-10 < x_1, x_2 < 10$. No jitter was included in this figure; when it was included the errors increased by a small amount. The sensory array consisted of 400 neurons with Gaussian tuning curves of width 2. The motor array had 100 neurons with cosine tuning curves of width π . The vector method of decoding was used and in the synaptic modification rule $k = 2$.

bott, 1994). The sensory inputs contain the same amount of noise for all combinations of x_1 and x_2 , thus making the low-amplitude tuning curves, when r is small, more susceptible to noise degradation.

Discussion

We have analyzed the interface where sensory information is transferred to a motor system to generate movement. By decoding we have interpreted spiking activity in the motor network as the command to generate movement to a particular location. This technique is powerful because it gives neuronal activity an unambiguous meaning that can be studied quantitatively. Information about the position of an external object was encoded by the activities of a modeled neural population and we focused on the crucial step where this information is transmitted to a motor array through a set of synapses. The core of this study relates to the characteristics of these synapses and the mechanisms by which they can be established.

We have derived a condition that guarantees proper alignment so that sensory input produces the desired motor output. This condition is just that the weight of the synapse from a presynaptic neuron of the sensory array to a postsynaptic neuron coding motor output depends only on the magnitude of the difference between their preferred locations. Synaptic weights in motor cortex with exactly this property have been inferred from correlation data (Georgopoulos et al., 1994). We found that appropriate synaptic weights can arise spontaneously during development or training if the motor system output provides random target location input to the sensory array and if correlation-based potentiation occurs on synapses from the sensory to the motor array. In contrast, network models based on learning rules

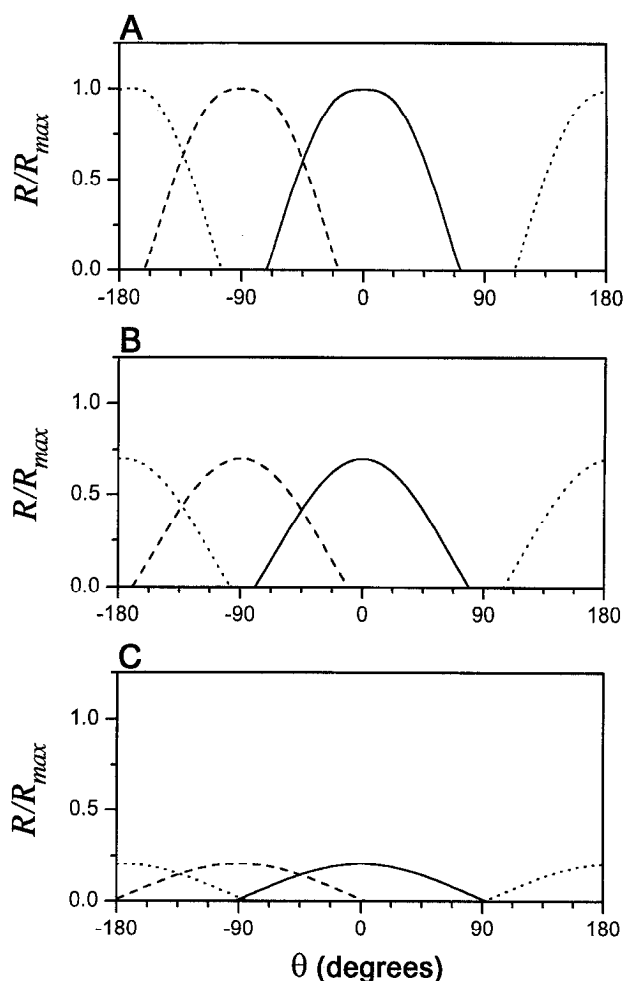


Figure 8. Evoked response tuning curves for neurons of the motor array. The average evoked firing rates of three neurons from the motor array of Figure 7 are shown plotted against the polar angle of the target location θ . Responses are shown for three different values of the length r of the target location vector. The result is a representation of the polar coordinates of the target with the polar angle tuning curves modulated as a function of radial length. A, $r = 8$; B, $r = 2$; C, $r = 0.5$.

such as back-propagation (Hertz et al., 1991) require an error signal to be computed and fed back throughout the network. This kind of learning is generally considered unlikely in biological contexts (Crick, 1989). The mechanism we have studied produces sensory-motor alignment without making use of either an error signal or any kind of external supervision. The only ingredients are random input to the sensory network that is correlated with the output of the motor network and a correlation-based synaptic modification rule. Randomly waving the arm around in front of the eyes and watching it is sufficient to align the sensory and motor networks. Studies have shown that vision of a limb is indeed required in motor development in both cats and monkeys (Hein, 1974; Held and Bauer, 1974), suggesting that this form of learning might be operating. In addition, a recent study has shown that human infants tend to move a viewed arm more than an arm that cannot be seen and that they require sight of an arm to counter deflecting forces (van der Meer et al., 1995).

For clarity, we have chosen to describe reaching tasks using the location of the target and the goal location of the resulting arm movement. It is also possible to use the required and actual

direction of arm movement in place of the target and goal location. These are just the target and goal location in hand-centered coordinates. Since both arm position and movement direction are simultaneously encoded in premotor and motor cortex (Schwartz et al., 1988; Caminiti et al., 1990, 1991), it should be possible for either of these two representations to be used. When random arm movements are observed, we have assumed that a representation of the present location of the arm in the motor array is correlated with a representation of target location in a sensory array. It is probably more appropriate to say that during learning the actual direction of random arm movements is correlated with their observed direction. This is equivalent to describing the learning process in hand-centered coordinates rather than in the body-centered coordinates that we have used. Similarly, during sensory-guided movement, we can think of the sensory representation of the target as evoking a motor response encoding the direction of movement toward the target rather than a goal location.

Noise is an important factor in the interpretation of our results. We have included large random variations in the neuronal responses and the alignment errors shown depend on the amount of noise included. We chose a rather high coefficient of variation in order to verify our results under worst case conditions. Without the random fluctuations the results shown would have been more accurate by at least two orders of magnitude and, in some cases, up to the accuracy of the numerical methods implemented. We have included examples of this in Figures 4, *A* and *B*, and 7 (dashed lines). This implies that the algorithms used to produce the sensory-motor alignment are very close to optimal so that noise limits the accuracy of the information transfer. The only exception is Figure 6, where the presence of additional variables simultaneously encoded by the sensory array essentially acts as an additional source of noise. In this case, the random fluctuations make less of a difference as the number of simultaneously encoded variables increases. Some of what we interpret as cortical noise may, in fact, be additional coded information of this sort that has not been correlated with any experimental manipulations.

The noise that we included in our simulations was uncorrelated between neurons. Experimental data suggest that, at least in some areas, neuronal noise is correlated (Gawne and Richmond, 1993; Shadlen and Newsome, 1994; Zohary et al., 1994). These results suggest that correlated noise effects become important for populations of about 100 neurons or more. Correlated noise will limit the $N^{-1/2}$ improvement of accuracy that we have seen for increasing numbers of coding neurons. Our results show that even with 100 neurons or less, fairly accurate transfer and coordinate transformations of target position can be realized. However, extremely high accuracy may require some additional mechanisms to decorrelate noise effects if correlated noise is indeed a widespread phenomenon.

Our model involved firing rates, not spikes. More recently we have found that similar results can be obtained using model neurons generating Poisson spike trains (Salinas and Abbott, unpublished observations). In very rapid tasks, there may only be time for a few spikes to be generated and thus only a rough approximation of firing rate can be conveyed by any single neuron. This uncertainty acts effectively as an additional source of uncorrelated noise and its impact is reduced by population coding. Shadlen and Newsome (1994) point out that, with about 100 neurons firing at normal rates, an accurate determination of

firing rate can be obtained from measurements of single interspike intervals.

The results presented provide an interesting view of how coordinate transformations are performed by neural networks. We focused on one particular coordinate transformation, that from retinal coordinates of target position to a head-centered, gaze-angle-invariant motor representation. However, the results are quite general. They apply whenever a motor network coding movement location is driven by a sensory array coding two or more quantities simultaneously. We have shown that synaptic corrections that arise spontaneously from correlation-based synaptic plasticity allow the motor network to extract vector sums, differences, and other linear or even nonlinear combinations from the sensory array. This provides a general mechanism for doing vector arithmetic and for making coordinate transformations like those from retinal to head-centered, head-centered to body-centered, or body-centered to arm-centered systems (Caminiti et al., 1990, 1991; Flanders et al., 1992; Kalaska and Crammond, 1992; Soechting and Flanders, 1992). The motor array we studied extracts the location of the target in head-centered coordinates. If this head-centered representation is encoded simultaneously with a representation of head position, a subsequent network can extract a body-centered representation of target location by the same mechanism (Andersen, 1989; Stein, 1992; Pouget and Sejnowski, 1994; Brotchie et al., 1995).

When we included a variable direction of gaze, we found that the receptive fields of neurons in the motor layer of our model shifted as a function of retinal position. We have studied the case of gaze angle and retinal position, but the same general mechanism can account for shifting responses in other modalities. For example, the preferred location vector of neurons in motor and premotor cortex can shift as a function of initial arm position (Caminiti et al., 1990, 1991). These shifts can, in principle, be accounted for by the mechanism we have studied provided that the appropriate simultaneous representation exists in arrays driving the neurons with shifting fields.

Although we have focused on information transfer between sensory and motor systems, our results can be applied to purely sensory cases as well. In particular, there are a number of examples of shifting receptive fields in a variety of sensory modalities. Neurons with shifting visual receptive fields have been recorded in parietal cortex (Duhamel et al., 1992). These dynamic shifts can be explained if the difference between intended direction of gaze and actual gaze direction is combined with retinal position by the mechanism we have discussed. The receptive fields of the driven array will then temporarily shift in response to an intended saccade before the visual image actually moves (Goldberg and Colby, 1992). Shifting auditory receptive fields have also been seen within area LIP (Stricanne et al., 1994). The receptive fields of certain neurons with visual responses in the premotor cortex of monkeys shift as a function of arm (Graziano et al., 1994) and head (Gross and Graziano, 1995) position; that is, their visual fields are in arm- and head-centered coordinates, respectively. Shifting receptive fields might at first appear quite mysterious since it may seem as if the pattern of visual inputs to a given neuron is being actively switched (Anderson and Van Essen, 1987; Olshausen et al., 1993). The model we have studied explains these shifts in terms of a gain modulation mechanism in the layers feeding the cells that actually display the shifting receptive fields (Zipser and Andersen, 1988; Andersen et al., 1993). In particular, the appearance of shifting receptive fields, even in parietal areas, should

not be viewed as contradictory to the observation of gain-modulated, nonshifting fields in area 7a. Rather, if this mechanism is operative their co-appearance is to be expected.

Premotor neurons with visual receptive fields in head-centered coordinates are particularly applicable to our study (Gross and Graziano, 1995). These cells respond to tactile stimulation on a region of the face and to visual stimuli near the same area of the face. Naturally, the tactile field moves when the head is rotated, but more remarkably, the visual field also rotates with the head but not with movement of the eyes (Gross and Graziano, 1995). This is exactly what we would expect from correlation-based synaptic modification. These neurons receive synaptic input from parietal neurons collectively and simultaneously encoding retinal location as well as head and eye orientation, so the necessary synaptic connections can arise spontaneously. In this case, the correlations driving the learning are not observed movements but rather correlations between visual and tactile responses. These fields would develop, for example, if visual stimuli in the region of the tactile field were repeatedly accompanied by tactile stimulation, as is likely. In this case, a postsynaptic premotor neuron is excited by the tactile stimulus while presynaptic parietal neurons are responding to the visual stimulus modulated by eye and head position. The tactile response will only be correlated with that combination of retinal, head, and eye position responses corresponding to the location of the visual stimulus relative to the face. As a result of correlation-based modification, synapses conveying this information will be strengthened selectively and will ultimately transfer it to the premotor neurons evoking their head-dependent, gaze-invariant visual responses. The fact that the network shown in Figure 4 could extract either a gaze-angle-dependent or gaze-angle-invariant representation has direct application to this case. We would predict that it might be possible to train these visual-tactile neurons to react to visual stimuli in laboratory or retinal coordinates rather than in head-centered coordinates if the correspondence between visual and tactile stimulation is appropriately modified over repeated trials.

It is tempting to think of neural networks as representing spatial information in particular coordinate systems. However, a number of researchers have pointed out that this may be inappropriate (Zipsper and Andersen, 1988; Andersen, 1989; Stein, 1992; Pouget and Sejnowski, 1994; Sanger, 1994) and our results support this latter viewpoint. If target location in retinal coordinates is represented simultaneously with information about gaze direction, any linear combination of these two quantities can be extracted by a downstream network. This means that the sensory array is really representing target location in a coordinate-invariant manner and arbitrary gaze-angle-dependent translations can easily be made. A network coding both arm position and movement direction can be considered to represent the goal location of the movement in a similarly coordinate-independent manner. We have also shown that a neuronal array can extract polar coordinates from another array responding to the Cartesian coordinates of a target location. This means that we should not consider the representation to be either Cartesian or polar. Again, it is better to think of these representations as coordinate independent. The advantage of such a scheme is clear: it allows the desired information, target location in a particular coordinate system relative to a particular point, to be extracted quite easily.

Our results have implications for mechanisms by which general invariant representations might arise (Fukushima, 1980;

Buhmann et al., 1990; Foldiak, 1991; Wallis et al., 1993). When correlation-based learning is used to establish synaptic connections, a network automatically selects the information that it needs to carry out its function. Information that is irrelevant to the output of a network will be eliminated from the synaptic input because it is not correlated with the activity of that network. As a result, a downstream network will automatically develop invariant responses to the combination of sensory information most correlated with its functional role. This is a powerful mechanism for channeling information effectively between different sensory and motor networks.

Appendix A

If a motor array is driven by a sensory array coding a target location x according to Equation 7 and if the synaptic weights satisfy Equation 15, the goal location of the movement z coded by the motor array will satisfy $z = x$. Average firing rates in the motor array can be determined by substituting the average rates in the sensory array (Eq. 1) into Equation 7 and ignoring noise terms:

$$R_i^m = \sum_j W_{ij} f(|a_j - x|). \quad (18)$$

According to the vector decoding method (Eq. 13),

$$z = \frac{\sum_i R_i^m c_i}{\sum_i R_i^m} = \frac{\sum_{ij} W_{ij} f(|a_j - x|) c_i}{\sum_{ij} W_{ij} f(|a_j - x|)}. \quad (19)$$

By replacing sums over neurons by integrals over their preferred locations and using Equation 15, we can rewrite this equation as

$$z = \frac{\int da dc W(|c - a|) f(|a - x|) c}{\int da dc W(|c - a|) f(|a - x|)}. \quad (20)$$

Shifting the integration variables $a \rightarrow a + x$ and $c \rightarrow c + x$ and noting that the integral

$$z = \int da dc W(|c - a|) f(|a|) c = 0 \quad (21)$$

(proved by making the substitutions $c \rightarrow -c$ and $a \rightarrow -a$) gives

$$z = \frac{\int da dc W(|c - a|) f(|a|) x}{\int da dc W(|c - a|) f(|a|)} = x. \quad (22)$$

Thus, synaptic weights satisfying Equation 15 will produce a motor response perfectly aligned to the location of the target.

Appendix B

The learning procedure leading to Equation 11 for the synaptic weights produces synaptic weights satisfying Equation 15 and thus aligns the response. To prove this we must show that the synaptic weights of Equation 11 depend only on the magnitude of the difference between the preferred locations for the pre- and postsynaptic neurons $|c_i - a_j|$. This is done by shifting the integration variable in Equation 11 $z \rightarrow z + a_j$, so that

$$W_{ij} = \int dz g(|c_i - a_j - z|) f(|z|) - k. \quad (23)$$

This is obviously function of $c_i - a_j$. The change of variables $z \rightarrow -z$ shows that it is in fact a function of $|c_i - a_j|$.

Appendix C

Synaptic weights satisfying Equation 16 evoke firing rates in the motor array, as given by Equation 7, that encode a goal location

$z = x + y$ if the sensory array encodes x and y simultaneously. On average in this case, ignoring noise terms,

$$R_i^m = \sum_j W_{ij} f(|a_j - x|, |b_j - y|), \quad (24)$$

and for the vector method of decoding,

$$z = \frac{\sum_i R_i^m c_i}{\sum_i R_i^m} = \frac{\sum_{ij} W_{ij} f(|a_j - x|, |b_j - y|) c_i}{\sum_{ij} W_{ij} f(|a_j - x|, |b_j - y|)}. \quad (25)$$

In integral form using Equation 16 this equation becomes

$$z = \frac{\int da db dc W(|c - a - b|) f(|a - x|, |b - y|) c}{\int da db dc W(|c - a|) f(|a - x|, |b - y|)}. \quad (26)$$

Shifting the integration variables $a \rightarrow a + x$, $b \rightarrow b + y$, and $c \rightarrow c + x + y$ and noting that the integral

$$z = \int da db dc W(|c - a - b|) f(|a|, |b|) c = 0 \quad (27)$$

(by making the substitutions $c \rightarrow -c$, $a \rightarrow -a$ and $b \rightarrow -b$) gives

$$z = \frac{\int da db dc W(|c - a - b|) f(|a|, |b|) (x + y)}{\int da db dc W(|c - a - b|) f(|a|, |b|)} = x + y. \quad (28)$$

Thus, the motor response z is equal to the linear combination of retinal and gaze directions $x + y$.

Appendix D

The expression (Eq. 12) for the synaptic weights arising from a correlation-based learning rule satisfies Equation 16 and thus produces a correctly transformed motor response. This requires that Equation 12 be a function of $|c_i - a_j - b_j|$. To show this we shift the integration variables in Equation 12 $y \rightarrow y + b_j$ and $z \rightarrow z + a_j + b_j$ and write

$$W_{ij} = \int dy dz g(|c_i - a_j - b_j - z|) f(|y - z|, |y|) - k. \quad (29)$$

To see that this is a function of $|c_i - a_j - b_j|$, make the changes $y \rightarrow -y$ and $z \rightarrow -z$.

References

Abbott LF (1994) Decoding neuronal firing and modeling neural networks. *Q Rev Biophys* 27:291–331.
 Abbott LF, Blum K (1995) Functional significance of long-term potentiation for sequence learning and prediction. *Cereb Cortex*, in press.
 Alexander GE, Crutcher MD (1990a) Preparation for movement: neural representations of intended direction in three motor areas of the monkey. *J Neurophysiol* 64:133–150.
 Alexander GE, Crutcher MD (1990b) Neural representation of the target (goal) of visually guided arm movements in three motor areas of the monkey. *J Neurophysiol* 64:164–178.
 Andersen RA (1989) Visual and eye movement functions of posterior parietal cortex. *Annu Rev Neurosci* 12:377–403.
 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, Snyder LH, Li C-S, Stricanne B (1993) Coordinate transformations in the representation of spatial information. *Curr Opin Biol* 3:171–176.
 Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. *Science* 230:450–458.
 Anderson CH, Van Essen DC (1987) Shifter circuits: a computational strategy for dynamic aspects of visual processing. *Proc Natl Acad Sci USA* 84:6297–6301.
 Biguer B, Donaldson IML, Hein A, Jeannerod M (1988) Neck muscle vibration modifies the representation of visual motion and direction in man. *Brain* 11:1405–1424.

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.
 Buhmann J, Lades M, von der Malsburg C (1990) Size and distortion invariant object recognition by hierarchical graph matching. In: Proceedings of the International Joint Conference on Neural Networks, pp 411–416. San Diego: IEEE.
 Bullock D, Grossberg S (1988) Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychol Rev* 95:49–90.
 Burnod Y, Grandguillaume P, Otto I, Ferraina S, Johnson PB, Caminiti R (1992) Visuomotor transformation underlying arm movements toward visual targets: a neural network model of cerebral cortical operations. *J Neurosci* 12:1435–1452.
 Caminiti R, Johnson PB, Urbano A (1990) Making arm movements within different parts of space: dynamical aspects on the primate motor cortex. *J Neurosci* 10:2039–2058.
 Caminiti R, Johnson PB, Galli C, Ferraina S, Burnod Y (1991) Making arm movements within different parts of space: the premotor and motor cortical representations of a coordinate system for reaching to visual targets. *J Neurosci* 11:1182–1197.
 Crick F (1989) The recent excitement about neural networks. *Nature* 337:129–132.
 Crutcher MD, Alexander GE (1990) Movement-related neuronal activity selectively coding either direction or muscle pattern in three motor areas of the monkey. *J Neurophysiol* 64:151–163.
 Duhamel J-R, Colby CL, Goldberg ME (1992) The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255:90–92.
 Flanders M, Helms-Tillery SI, Soechting JF (1992) Early stages in a sensorimotor transformation. *Behav Brain Sci* 15:309–362.
 Foldiak P (1991) Learning invariance from transformed sequences. *Neural Comput* 3:194–200.
 Fortier PA, Kalaska JF, Smith AM (1989) Cerebellar neuronal activity related to whole-arm reaching movements in the monkey. *J Neurophysiol* 62:198–211.
 Fukushima K (1980) Neocognitron: a self-organized neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol Cybern* 36:193–202.
 Gaudio P, Grossberg S (1991) Vector associative maps: unsupervised real-time error based learning and control of movement trajectories. *Neural Networks* 4:147–183.
 Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? *J Neurosci* 13:2758–2771.
 Georgopoulos AP, Schwartz A, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419.
 Georgopoulos AP, Kettner RE, Schwartz A (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *J Neurosci* 8:2928–2937.
 Georgopoulos AP, Taira M, Lukashin A (1993) Cognitive neurophysiology of the motor cortex. *Science* 260:47–52.
 Glickstein M (1990) Brain pathways in the visual guidance of movement. In: *Brain circuits and functions of the mind* (Trevarthen C, ed), pp 157–167. Cambridge, MA: Cambridge UP.
 Goldberg ME, Colby CL (1992) Oculomotor control and spatial processing. *Curr Opin Neurobiol* 2:198–202.
 Graziano MSA, Yap GS, Gross CG (1994) Coding of visual space by premotor neurons. *Science* 266:1054–1057.
 Gross CG, Graziano MSA (1995) Multiple representations of space in the brain. *Neuroscientist* 1:43–50.
 Hebb DO (1949) *The organization of behavior: a neuropsychological theory*. New York: Wiley.
 Hein A (1974) Prerequisite for development of visually guided reaching in kittens. *Brain Res* 71:259.
 Held R, Bauer JA (1974) Development of sensorially-guided reaching in infant monkeys. *Brain Res* 71:265.
 Hertz J, Krogh A, Palmer RG (1991) *Introduction to the theory of neural computation*. New York: Addison-Wesley.
 Jeannerod M, ed (1990) *Attention and performance*, Vol 13, Motor representation and control. Hillsdale, NJ: Erlbaum.
 Kalaska JF, Crammond DJ (1992) Cerebral cortical mechanisms of reaching movements. *Science* 255:1517–1522.
 Kawato M, Isobe M, Maeda Y, Suzuki R (1988) Coordinate transfor-

- mations and learning control for visually guided voluntary movements with iteration: a Newton-like method in function space. *Biol Cybern* 59:161–177.
- Kettner R, Marcario J, Port N (1993) A neural network model of cortical activity during reaching. *J Cognit Neurosci* 5:14–33.
- Knudsen E (1982) Auditory and visual maps of space in the optic tectum of the owl. *J Neurosci* 2:1177–1194.
- Knudsen E, Konishi M (1978) A neural map of auditory space in the owl. *Science* 200:795–797.
- Kristan WB (1982) Sensory and motor neurons responsible for local bending responses in leeches. *J Exp Biol* 96:161–180.
- Kuperstein M (1988a) Neural model of adaptive hand-eye coordination for single postures. *Science* 239:1208–1311.
- Kuperstein M (1988b) An adaptive model for mapping invariant target position. *Behav Neurosci* 102:148–162.
- Lee C, Rohrer WH, Sparks DL (1988) Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature* 332:357–360.
- Lockery SR, Kristan WB (1990a) Distributed processing of sensory information in the leech. I. Input-output relations of the local bending reflex. *J Neurosci* 10:1811–1815.
- Lockery SR, Kristan WB (1990b) Distributed processing of sensory information in the leech. II. Identification of interneurons contributing to the local bending reflex. *J Neurosci* 10:1816–1829.
- Lukashin AV, Georgopoulos AP (1993) A dynamical neural network model for motor cortical activity during movement: population coding of movement trajectories. *Biol Cybern* 69:517–524.
- Mussa-Ivaldi FA, Giszter SF (1992) Vector-field approximation: a computational paradigm for motor control and learning. *Biol Cybern* 67:491–500.
- Olshausen BA, Anderson CH, Van Essen DC (1993) A neurobiological model of visual attention and invariant pattern recognition based on dynamical routing of information. *J Neurosci* 13:4700–4719.
- Pouget A, Sejnowski TJ (1994) A neural model of the cortical representation of egocentric distance. *Cereb Cortex* 4:314–329.
- Press WH, Flannery BP, Teukolsky SA, Vetterling WT (1992) *Numerical recipes in C*. New York: Cambridge UP.
- Redish AD, Touretzky DS (1994) The reaching task: evidence for vector arithmetic in the motor system. *Biol Cybern* 71:307–317.
- Roll R, Velary JL, Roll JP (1991) Eye and neck proprioceptive messages contribute to the spatial coding of retinal input in visually oriented activities. *Exp Brain Res* 85:423–431.
- Salinas E, Abbott LF (1994) Vector reconstruction from firing rates. *J Comput Neurosci* 1:89–107.
- Sanger TD (1994) Theoretical considerations for the analysis of population coding in motor cortex. *Neural Comput* 6:29–37.
- Schwartz A, Kettner RE, Georgopoulos AP (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *J Neurosci* 8:2913–2927.
- Sejnowski TJ (1977) Storing covariance with nonlinearly interacting neurons. *J Math Biol* 4:303–321.
- Shadlen MN, Newsome WT (1994) Noise, neural codes and cortical organization. *Curr Opin Neurobiol* 4:569–579.
- Soechting JF, Flanders M (1992) Moving in 3-dimensional space: frames of reference, vectors and coordinate systems. *Annu Rev Neurosci* 15:167–191.
- Softky WR, Koch C (1992) Cortical cells should spike regularly but do not. *Neural Comput* 4:643–646.
- Softky WR, Koch C (1994) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350.
- Stein J (1992) Representations of egocentric space in posterior parietal cortex. *Brain and Behav Sci* 15:691–703.
- Stricanne B, Xing J, Mazzoni P, Andersen RA (1994) Response of LIP neurons to auditory targets for saccadic eye movements: a distributed coding for sensorimotor transformation. *Soc Neurosci Abstr* 20:143.
- Takahashi TT (1989) The neural coding of auditory space. *J Exp Biol* 146:307–322.
- van der Meer ALH, van der Weel FR, Lee DN (1995) The functional significance of arm movements in neonates. *Science* 267:693–695.
- van Gisbergen JAM, van Opstal AJ, Tax AMM (1987) Collicular ensemble coding of saccades based on vector summation. *Neuroscience* 21:541–555.
- Wallis G, Rolls E, Foldiak P (1993) Learning invariant responses to the natural transformations of objects. *Int Joint Conf Neural Networks* 2:1087–1090.
- 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.