Rank-Order-Selective Neurons Form a Temporal Basis Set for the Generation of Motor Sequences

Emilio Salinas

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

Many behaviors are composed of a series of elementary motor actions that must occur in a specific order, but the neuronal mechanisms by which such motor sequences are generated are poorly understood. In particular, if a sequence consists of a few motor actions, a primate can learn to replicate it from memory after practicing it for just a few trials. How do the motor and premotor areas of the brain assemble motor sequences so fast? The network model presented here reveals part of the solution to this problem. The model is based on experiments showing that, during the performance of motor sequences, some cortical neurons are always activated at specific times, regardless of which motor action is being executed. In the model, a population of such rank-order-selective (ROS) cells drives a layer of downstream motor neurons so that these generate specific movements at different times in different sequences. A key ingredient of the model is that the amplitude of the ROS responses must be modulated by sequence identity. Because of this modulation, which is consistent with experimental reports, the network is able not only to produce multiple sequences accurately but also to learn a new sequence with minimal changes in connectivity. The ROS neurons modulated by sequence identity thus serve as a basis set for constructing arbitrary sequences of motor responses downstream. The underlying mechanism is analogous to the mechanism described in parietal areas for generating coordinate transformations in the spatial domain.

Introduction

Many sophisticated behaviors, from playing a Bach prelude on the piano to tying a knot, are organized as strings of motor actions that must be performed in the proper order and at precisely the right times. Many simpler behaviors, however, such as eating a banana or climbing a tree, also involve the sequencing of multiple, separate movements. Thus, motor sequencing is a fundamental function of the brain. Not surprisingly, a variety of cortical and subcortical structures are involved in the learning, storage, and execution of movement sequences (Tanji, 2001; Ashe et al., 2006). In general, however, evidence from lesion, imaging, and neurophysiological studies indicates that, in primates, there are three cortical structures that are critical specifically for planning and executing sequences of discrete movements over time: the supplementary motor area (SMA), the pre-SMA, and the supplementary eye field (SEF) (Tanji, 2001; Nachev et al., 2008). A key observation is that, when the SMA is pharmacologically inactivated, monkeys become incapable of making sequential movements from memory, although they can still perform the individual movement components (Gerloff et al., 1997; Shima and Tanji, 1998).

Single-neuron recordings in behaving monkeys trained to perform instructed motor sequences have provided a wealth of information about the neural basis of this capacity. In particular, studies by Tanji and collaborators have documented the activity

of so-called rank-order-selective (ROS) neurons, which are abundant precisely in the pre-SMA, SMA, and SEFs and are also found in the basal ganglia, in which they are known as phase-selective neurons (Kermadi and Joseph, 1995; Mushiake and Strick, 1995).

ROS cells fire during sequences of arm movements (Mushiake et al., 1991; Clower and Alexander, 1998; Shima and Tanji, 1998, 2000) or saccades (Isoda and Tanji, 2003, 2004; Averbeck et al., 2006; Berdyeva and Olson, 2009). Their defining characteristic is that they are active during specific parts of a motor sequence. For instance, suppose sequence 1 is pull–push–turn, meaning that the monkey has to pull a key, then push it, and then turn it, sequence 2 is pull–turn–pull, and so on for other combinations. An ROS neuron could be active, say, during the transition between the second and third movements in all sequences, regardless of the actual movements performed at those points. That is, ROS cells are active at fixed time intervals during a multipart motor action, and this activity is most consistent with a preference for a particular serial position in a sequence (Berdyeva and Olson, 2009).

In addition, there is another feature of these cells that is suggested by the experimental reports and that, according to the model presented here, is also crucial for their function: their overall response amplitude depends on sequence identity. Thus, over a variety of sequences, an ROS unit is always activated during the same time period but with varying intensities.

This study shows that ROS neurons that are modulated by sequence identity in essence solve the problem of assembling arbitrary sequences of motor actions. The results are based on theoretical calculations and computer simulations of a model network in which such cells serve to construct many motor se-

Received Jan. 12, 2009; revised Feb. 14, 2009; accepted Feb. 19, 2009.

This work was supported in part by National Institute of Neurological Disorders and Stroke Grant NS044894 (E.S.). 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@wfubmc.edu.

DOI:10.1523/JNEUROSCI.0164-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/294369-12\$15.00/0

quences composed of a set of discrete movements arranged in different orders. Interestingly, the underlying mechanism is similar to the nonlinear mechanism thought to be the basis for the computation of coordinate transformations in the visual system, except that here it is applied to the time domain.

Materials and Methods

All simulations were performed using Matlab (MathWorks). The code is available from the author on request.

Network architecture and simulated task. The model network has two layers. The first contains ROS neurons, and these drive a second layer of output or motor neurons through a set of synaptic connections. In each simulated trial, the ROS units are activated and their responses drive the motor units, which produce a sequence of three motor actions. The individual motor actions are denoted as A, B, and C; they last 1 s and are separated from each other by additional blank or no-movement intervals also lasting 1 s. Each three-movement sequence has the structure -X-X-X- where X stands for one of the three movements and each dash indicates a blank interval. Thus, each simulated trial lasts a total of 7 s. This structure is not essential; it was chosen for simplicity and because it approximates the structure of the tasks used by Tanji and colleagues and by other groups (Shima and Tanji, 1998, 2000; Lu et al., 2002; Isoda and Tanji, 2003, 2004; Averbach et al., 2006). It is assumed that a cue at a beginning of a trial indicates which sequence is to be performed (from memory). The cue itself is not simulated, but its effect on the ROS cells is.

Rank-order-selective responses. The responses of the ROS neurons depend on time and on the identity of the sequence to be performed. The symbol R_{jqt} identifies the instantaneous firing rate of ROS unit j at time step t during sequence q . The neurophysiological results suggest that the time dependence of an ROS neuron is the same for all sequences and that sequence identity modulates the overall amplitude of its response. Therefore, the firing rate of ROS neuron j is modeled as

$$r_{jqt} = r_{\min} + r_{\max} g_{jq} f_{jt} \quad (1)$$

$$R_{jqt} = r_{jqt} + \varepsilon, \quad (2)$$

where r_{jqt} is the mean firing rate, averaged over trials, and ε represents random noise. In addition, $r_{\min} = 2$ and $r_{\max} = 33$ spikes/s are constant terms, g_{jq} is the gain of neuron j during sequence q , and f_{jt} is the temporal profile of the neuron, i.e., its normalized response as a function of time in any sequence. In the simulations, the gain factors g_{jq} for each model neuron vary randomly between $g_{\min}$ and 1 across sequences, which means that the maximum possible response suppression from one sequence to another, expressed as a percentage, is equal to $(1 - g_{\min}) \times 100\%$. The temporal profiles f_{jt} are chosen so that each neuron is active for ~ 1000 ms, either during one of the movement periods or during one of the preparatory periods that precede them.

To include neuronal variability, a noise term ε is added to the response of every ROS unit (Eq. 2). A different ε is drawn from a Gaussian distribution for each cell at every point in time. The variability is modeled as Poisson like, meaning that the variance in the firing rate of unit j is equal to α times the mean rate of unit j . Therefore, because the mean firing rate of unit j at the current time is r_{jqt} , the mean and variance of ε are

$$\langle \varepsilon \rangle = 0 \quad (3)$$

$$\langle \varepsilon^2 \rangle = \alpha r_{jqt} \quad (4)$$

where the angle brackets indicate an average over trials, or repetitions. These random fluctuations are uncorrelated across cells.

In Equation 1, the temporal and sequence dependencies are combined multiplicatively. Although possible (Salinas and Abbott, 1996; Peña and Konishi, 2001; Sohn and Lee, 2007), an exact multiplication is not necessary; the obligatory condition is that the dependencies must be combined nonlinearly (Pouget and Sejnowski, 1997; Salinas, 2004a,b). This means that linear combinations should be avoided. For instance, when the mean firing rate of ROS neuron j during sequence q is described by an additive combination of time- and sequence-dependent terms, as in

$$r_{jqt} = r_{\min} + r_{\max} (g_{qt} + f_{qt}), \quad (5)$$

the model fails completely. As long as the interaction remains nonlinear, however, the results should vary very little with respect to those obtained with Equation 1 (Salinas, 2004a,b).

Motor responses. The symbol m_{kqt} stands for the firing rate of motor unit k during time point t of sequence q and is calculated through a weighted sum of ROS responses,

$$m_{kqt} = \sum_j w_{kj} R_{jqt}, \quad (6)$$

where w_{kj} represents the synaptic connection from ROS neuron j to output neuron k . This expression is used when the ROS neurons drive the motor neurons, i.e., after the synaptic connections have been established. However, there is also an intended or desired motor response for each output neuron. It is denoted as M_{kqt} and is used only for setting the synaptic connections. This is as follows.

In the simulations, each motor unit is supposed to contribute to the generation of one of the three possible movements, A, B, or C. Initially, the desired response of each motor neuron is simply a square function of time that is on only during the associated movement period or during the preceding preparatory period and is off otherwise. For instance, if motor neuron k is associated with the preparation of movement A, then during the sequence ABC (sequence 1), its desired motor response would be

$$M_{k1t} = \begin{cases} 1 & \text{if } 0 \leq t < 1000 \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

because the preparatory period before A starts at time 0 in this sequence. In contrast, during the sequence BAC (sequence 2), the desired motor response for the same cell would be

$$M_{k2t} = \begin{cases} 1 & \text{if } 2000 \leq t < 3000 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

because the preparatory period before A now starts 2 s after the beginning of the trial, and so on for other sequences. For clarity, the t index in these expressions runs in 1 ms increments but, in the simulations, every increment was equal to a 10 ms step. Finally, all the desired motor firing rates are smoothed with a Gaussian filter with an SD of 50 ms, and the results are scaled and added to a background firing rate (for examples of the final shapes of the desired motor profiles, see Figure 3c). Equal numbers of motor units are assigned to each of the three movements.

Synaptic weights. The connections from the ROS to the motor neurons are chosen so that they minimize the average squared difference between the desired and driven responses; that is, they must make the error

$$E = \sum_{k,q,t} \langle \phi_q (M_{kqt} - \sum_j w_{kj} R_{jqt})^2 \rangle \quad (9)$$

as small as possible, where the angle brackets indicate an average over trials. Such minimization is a standard procedure for linear networks (Haykin, 1999; Salinas, 2004a,b). In this expression, the coefficients ϕ_q have been included to regulate the relative importance of different sequences, and thus the relative accuracy with which they are generated by the network. These importance coefficients satisfy the constraint $\sum_q \phi_q = 1$, where the sum is over all sequences. Except when explicitly indicated, all sequences have the same importance, so $\phi_q = 1/N_Q$ for all q , where N_Q is the total number of sequences.

The matrix of optimal connections $\mathbf{w}$ that minimizes the expression above is the one that satisfies the equation

$$\mathbf{wC} = \mathbf{L} \quad (10)$$

where

$$C_{jk} = \sum_{q,t} \phi_q [r_{jqt} r_{kqt} + \alpha \delta_{jk} r_{jqt}] \quad (11)$$

$$L_{jk} = \sum_{q,t} \phi_q M_{jqt} r_{kqt} \quad (12)$$

The second term in Equation 11 is attributable to the variability (noise) in the ROS responses. Its overall strength is determined by α , where $\alpha = 0$ means that there is no noise and $\alpha = 1$ means that the noise follows Poisson statistics. In this term, $\delta_{jk} = 1$ if $j = k$ and $\delta_{jk} = 0$ if $j \neq k$.

The solution to Equation 10 can be obtained by calculating the inverse (or the pseudo-inverse) of the correlation matrix C , in which case $\mathbf{w} = LC^{-1}$. Alternatively, in Matlab, the system of equations can be solved more efficiently through the slash operator, such that $\mathbf{w} = L/C$.

In selected simulations, the effect of random synaptic corruption is investigated by deleting some of the synapses in $\mathbf{w}$. This is controlled by the parameter P_w , which is the probability that any element w_{ij} is set to zero after training.

Three measures of network performance. Once the synaptic weights $\mathbf{w}$ are obtained, the motor responses driven by the ROS neurons are computed using Equation 6. Then, the actual differences between the desired and driven motor responses provide an indication of the accuracy of the network. The root-mean-square (RMS) error

$$E_{\text{RMS}} = \left[\frac{1}{N} \sum_{k,q,t} \langle (M_{kqt} - m_{kqt})^2 \rangle \right]^{1/2} \quad (13)$$

is used to quantify the magnitude of these differences, where N is equal to the number of motor neurons, times the number of sequences, times the number of time points per sequence.

Another way to quantify the performance of the network is to decode the movement that is being generated at each point in time and compare it with the intended movement at those same points. This is done with networks that include six motor neurons, one that fires before movement A, another that fires during movement A, another that fires before movement B, and so on. At each time point, the motor neuron with the maximum firing rate is identified. If this neuron is associated, say, with movement A and the desired movement at that time is indeed A, then that time point is scored as correct. Vice versa, if the unit with the highest rate is associated with a movement different from the desired one at that point, then that time point is scored as an error. In this way, by averaging over all sequences and multiple time points and trials, the probability of encoding an incorrect movement, P_m , is calculated. Time points near the borders between blank and movement intervals, during which the motor firing rates rise or fall, are excluded from the calculation.

The probability P_m quantifies a movement error that is brief, occurring over a single time step (of 10 ms in the simulations). To quantify movement errors that persist over the timescale of a whole movement, a similar procedure is implemented that scores each of the six movement periods as either correct or incorrect. For instance, suppose the first movement is A; the encoded movement is scored as correct if in at least one-half of the time points of the first movement period the neuron with the maximum firing rate is the one associated with A. Conversely, the encoded movement is scored as incorrect if in more than one-half of the time points during this movement period the neuron with the maximum firing rate is other than the one associated with A. By averaging over all sequences and time periods, and over trials, the probability of encoding an incorrect movement, P_M , is calculated. This quantity indicates how often the network will generate a wrong movement signal that lasts throughout one of the 1 s intervals.

Results

This study analyzes the functional relationship between two types of neurons, ROS neurons and motor neurons.

Shima and Tanji (2000) cataloged and described various subtypes of ROS neurons according to their temporal and sequence selectivities. The general class of ROS neurons considered here contains those neurons that they identified as ROS as well as other types that they discussed separately. Specifically, here, a neuron is considered as ROS if its temporal profile of activity remains approximately constant for all sequences, even if its overall response amplitude does change across sequences. This

Figure 1. Modulated ROS neurons can generate multiple motor sequences. The model network consisted of 420 ROS neurons driving, through synaptic connections, a population of 60 motor neurons. It was trained to produce six sequences of three movements each. The possible movements were A, B, and C. Colored arrays show the firing rates of the two populations. The x-axis is time, and each row on the y-axis is one neuron; for clarity, only a subset of neurons is shown. Sequence identity is indicated at the top. Horizontal bars mark the 1 s movement periods. Each ROS unit was active during a fixed time period in all sequences, but its response amplitude varied as a function of sequence identity. The shown responses are the mean firing rates expected after averaging over infinitely many trials. Here, $g_{\text{min}} = 0.4$ and $\alpha = 1$. Firing rate units are spikes per second.

definition is consistent with their proposed functional role, as shown below.

Practically all areas in which ROS cells are found contain a second type of cell that fires most intensely during the execution of specific movements (Tanji and Shima, 1994; Shima and Tanji, 2000; Isoda and Tanji, 2003). For simplicity, here these neurons are referred to as motor units. They show little or no dependence on sequence identity or on movement order within a sequence. For example, a motor neuron selective for “turn” would be active whenever a turn of the key is executed. Also, no distinction is made between neurons that fire predominantly before the actual movement (preparatory neurons) or during its execution. Here, all these movement-locked cells are considered the outputs of the circuit.

Motor activity driven by ROS activity

The model network developed here has a simple two-layer architecture in which a population of ROS neurons modulated by sequence identity drives a set of motor neurons downstream, which generate movements. The two layers are connected through a set of synaptic weights $\mathbf{w}$, which are determined according to the identity of the sequences that the network is required to store (see Materials and Methods). The behavior of the network is first described with an example in which six sequences of three movements are performed. Each sequence is assembled by combining the movements A, B, and C. The responses of the ROS and motor populations during each of the six sequences are shown in Figure 1, which plots the firing rate of the cells, encoded

by color, as a function of time.

Each ROS unit has a fixed temporal profile of activity, so it fires intensely during a particular time interval in any sequence. For instance, the model cells at the top left corner of the ROS arrays always fire during the first preparatory period. Their maximum response amplitude, or gain, however, changes as a function of sequence identity. For each model cell in this example, the six amplitude factors varied between 0.4 and 1 and were assigned randomly across sequences. This modulation can be seen by comparing the color intensity of a given cell across different sequences. The model ROS population includes a variety of temporal profiles, as reported previously (Shima and Tanji, 2000). Thus, their various segment or rank-order preferences tile the full duration of a sequence.

The motor neurons, conversely, behave in the opposite way: they fire at different points in time, depending on the sequence, but do so with the same intensity in all of them (if the corresponding movement is part of the sequence). Each motor unit fires in association with one of the movements (A, B, or C), either during its execution or during the preceding preparatory interval. For example, the units at the top of the motor arrays in Figure 1 always fire during the preparatory period leading to movement A.

What the model network achieves is to produce responses that can occur at any point in time out of responses that occur at fixed times. Although it is not obvious, the key for this is the sequence selectivity of the driving ROS cells.

Figure 2 plots the firing rates of three ROS neurons (Fig. 2a) and three motor neurons (Fig. 2b) as functions of time during the six sequences. These responses are from the same simulation as in Figure 1, but the changes in gain are more evident in this format. Again, the key observation is that (1) the response of each ROS neuron occurs during a particular time interval but with different intensities across sequences, whereas (2) the response of each motor neuron has a constant amplitude but occurs at a different time interval in each sequence, depending on the time at which the corresponding motor movement needs to be executed.

For simplicity and to mimic typical experimental situations, three types of movement and three movements per sequence were used in these examples, but in principle, the mechanism works for any number of movements and sequences, as long as the ROS population contains the necessary combinations of temporal and sequence selectivities (see Appendix).

Network capacity and accuracy

A crucial question for this type of network is how its accuracy relates to its size and to the number of motor sequences it must generate. Here, “size” really refers to the number of ROS neurons, because the accuracy of the simulated motor responses does not

Figure 2. Firing rates as functions of time for six model neurons from the same simulation as in Figure 1. Each set of axes corresponds to one of the six sequences, as indicated above each graph. Gray rectangles mark the three movement periods per sequence. *a*, Responses of three ROS neurons, which are active during the first movement period (dark blue), during the second preparatory period (magenta), and during the third preparatory period (light blue). Note the variations in amplitude. Maximum suppression is $\sim 51\%$ for these cells. *b*, Responses of three motor neurons, which fire during movement A (dark blue), during the preparatory period preceding movement B (magenta), and during the preparatory period preceding movement C (light blue). Firing rate units are spikes per second.

depend on the number of motor neurons included in the network. To investigate the capacity of the network, motor accuracy will be quantified as a function of network size under various conditions. Two different measures of accuracy will be considered. The first one is the RMS difference between driven and desired motor firing rates, E_{RMS} (Eq. 13).

The results in Figures 1 and 2 were obtained after training the network to produce a series of desired motor responses (see Materials and Methods). Here, “training” means that, given the firing rates of the ROS neurons modulated by sequence identity, optimal synaptic connections are found such that the motor responses driven by those ROS cells approximate the desired responses as accurately as possible. The driven motor responses in Figures 1 and 2 are very close to those used during training but are not exactly the same. The typical difference between them is E_{RMS} . By means of this quantity, it will become clear that there are three potential sources of error in the network: the number of stored sequences relative to network size, the temporal structure of the ROS responses, and the trial-to-trial variability, or noise.

First consider an idealized case in which (1) there is no noise

Figure 3. Comparison between desired and driven motor responses in networks of different sizes and with different amounts of variability. **a**, RMS error, E_{RMS} , as a function of the number of ROS units in networks trained to perform six sequences of three movements each. Each data point is an average over 50 networks. Symbols indicate different simulation conditions. Circles, Without noise ($\alpha = 0$) and without variability across ROS temporal profiles. Squares, Without noise ($\alpha = 0$) and with variability across ROS temporal profiles. Triangles and diamonds, With noise ($\alpha = 1$) and with variability across ROS temporal profiles. Shown error values are for driven responses in single trials (diamonds) and for mean driven responses expected after averaging over infinitely many trials (triangles). **b–g**, Driven responses of two motor neurons as functions of time (blue and magenta traces) during a sequence. Thin black lines are the corresponding desired responses. Each panel is from a different network and corresponds to one of the red points in **a**, as labeled. The green triangle in **a** indicates the network used in Figures 1 and 2.

or response variability ($\alpha = 0$), and (2) all neurons, motor and ROS, have exactly the same temporal activation profiles. Thus, in contrast to Figure 2a, where each of the three ROS units has a slightly different activation width, onset time (relative to the start of the relevant period), and activation shape, now the time courses of all neurons have exactly the same shape. In this case, motor accuracy is determined primarily by the number of ROS neurons relative to the number of stored sequences. In particular, in theory, the minimum number of ROS neurons needed to store N_Q sequences of N_S steps or time periods each is $N_{\text{ROS}} = N_Q \times N_S$ (see Appendix). The circles in Figure 3a show the RMS error E_{RMS} as a function of N_{ROS} for networks that generate six sequences throughout seven time periods, as in Figures 1 and 2. The curve falls very quickly until it exceeds 42 ROS neurons, the critical number; beyond that, the residual error is very small. Figure 3, **b** and **c**, shows examples of driven and desired motor responses for this type of simplified network. With fewer than 42 ROS units (Fig. 3b, $N_{\text{ROS}} = 28$), the two driven motor responses shown are completely wrong during the third preparatory period: the ma-

genta neuron fires much less than it should, and the blue neuron fires much more than it should; the wrong movement is encoded in this period, although there is no noise in the system. In contrast, with >42 ROS units (Fig. 3c; $N_{\text{ROS}} = 91$), the driven motor responses are indistinguishable from the desired ones.

Next, consider a situation in which (1) there is still no noise but (2) the temporal profiles of the ROS units are not identical. In this case, each ROS unit is active for a slightly different amount of time (1000 ± 160 ms), its onset of activity is within 20 ms of the start of the corresponding period, and its profile is not flat at the top. Examples are shown in Figure 2a. Meanwhile, the desired motor responses are exactly the same as before. Now a network with 91 ROS neurons has an RMS error of approximately $E_{\text{RMS}} \approx 1.3$ spikes/s, and the distinction between networks with more or fewer than 42 ROS units is strongly blurred (squares in Fig. 3a). As shown in Figure 3d, although the driven motor responses are activated at the right times and with intensities that are on average very close to the correct levels, they are not smooth. There are many small deviations from the desired responses, which are attributable to the discrepancy between the temporal profiles of the ROS units and the profiles of the desired motor responses.

Finally, consider the more general case in which (1) there is response noise ($\alpha = 1$), and (2) there is variability across the ROS temporal profiles. Response noise, as one might expect, always increases the RMS error. For example, in a relatively small network ($N_{\text{ROS}} = 91$), with Poisson variability the error goes from ~ 1.3 to ~ 3.9 spikes/s (Fig. 3, compare **d**, **e**). The colored traces in Figure 3e illustrate the motor responses of two model cells in a single trial under these conditions. Comparing these traces with those in Figure 3g shows that the resulting single-trial fluctuations decrease substantially when the network grows to $N_{\text{ROS}} = 539$ neurons. The same thing can be seen following the diamonds in Figure 3a. Last, when there is response variability, as in real recordings, it is standard practice to average neuronal responses over multiple trials. The triangles in Figure 3a show the RMS error between the desired motor responses and the mean driven responses in the limit when these are averaged over infinitely many trials. The effect of trial averaging can also be seen by comparing Figure 3, **g** and **f**.

In summary, to reproduce a set of desired motor sequences accurately, a network must have at least a minimum number of modulated ROS neurons, which depends on the number and length of the sequences that must be generated. Additional neurons are necessary, however, to compensate for trial-to-trial random fluctuations and for differences between the temporal activation profiles of the ROS neurons and the desired motor activation profiles.

Other measures of motor accuracy

Another way to quantify the accuracy of the network is to compare the movement that it is supposed to produce at each point in time (A, B, or C) with the movement that it actually encodes. This is done with the quantities P_m and P_M (see Materials and Methods). P_m is the probability that the output layer of the network encodes the wrong movement at any time point during a simulated sequence. A coding error of this type applies to a single time step (of 10 ms), so it is a brief event. Figure 4a shows that this probability decreases steadily as the size of the network increases. The data along the thin black line are for networks trained to perform six sequences of three movements, as in Figure 3, whereas the data along the gray line are for networks trained to perform 18 sequences of three movements. The P_m values rise in

Figure 4. Probability of encoding an incorrect movement as a function of the number of ROS neurons in a network. **a**, The thin black line (exponential fit) and the dots over it (simulation results) correspond to the probability P_m of encoding an incorrect movement in one time step (of 10 ms) for networks that store $N_Q = 6$ sequences of three movements each. The gray line (exponential fit) and the dots over it (simulation results) indicate the probability P_m for networks that store $N_Q = 18$ sequences of three movements each. Other parameters are $g_{\min} = 0.4$ and $P_w = 0$. **b**, As in **a** but for the probability P_M of encoding an incorrect movement throughout one of the 1 s periods. Continuous lines are Gaussian fits. **c**, Probability P_m as a function of network size for $g_{\min} = 0$ (orange line) and $g_{\min} = 0.85$ (purple line). Both lines are exponential fits. For comparison, the thin black line is the same as in **a**. Other parameters are $N_Q = 6$ and $P_w = 0$. **d**, As in **c** but for the probability P_M . Continuous lines are Gaussian fits. The thin black line is the same as in **b**. **e**, Probability P_m as a function of network size for $P_w = 0.05$ (orange line) and $P_w = 0.25$ (purple line). Both lines are exponential fits. The thin black line is the same as in **a**. Other parameters are $N_Q = 6$ and $g_{\min} = 0.4$. **f**, As in **e** but for the probability P_M . Continuous lines are exponential fits. The thin black line is the same as in **b**. For all curves, $\alpha = 1$.

the latter case because, when the network has to store a larger number of sequences, each one is reproduced less accurately.

A similar trend is seen with P_M (Fig. 4*b*), which is the probability that the output layer of the network encodes the wrong movement throughout one of the movement or preparatory intervals. A coding error of this type occurs when a wrong movement is signaled during $>50\%$ of a movement or preparatory interval, so intuitively it corresponds more closely to the actual production of a wrong movement. The two lines and sets of data points in Figure 4*b* plot P_M versus network size for the same simulations as in Figure 4*a*.

Like E_{RMS} , the two probabilities for encoding an incorrect movement quantify the accuracy of the network, but they are much less sensitive than E_{RMS} to small differences between driven and desired responses. For instance, $P_m = 0$ for the network illustrated in Figure 3*d*, which has an RMS error above 1, and P_m is only 0.085 for the network illustrated in Figure 3*e*, which has an RMS error above 4. Furthermore, the curves for P_M fall much more sharply than those for P_m , because each of the long error events that count toward P_M requires at least 40 of the brief error events that count toward P_m . These differences are summarized

by the following observation: with noise, as in Figure 3*a* (diamonds and triangles) and Figure 4, *a* and *b*, E_{RMS} falls as $1/\sqrt{N_{\text{ROS}}}$, P_m falls exponentially with the number of ROS neurons ($\sim \exp(-N_{\text{ROS}}/\nu)$), and P_M falls even faster, with a Gaussian tail ($\sim \exp(-N_{\text{ROS}}^2/\sigma^2)$).

Network robustness

Four factors that affect network performance have been discussed: the size of the network, the number of stored sequences, the mismatch between ROS and motor temporal profiles, and response variability. There are two more components of the model that also determine performance: (1) the strength of the modulation by sequence identity, and (2) the precision of the synaptic weights.

The effect of modulation strength can be seen by varying the parameter $g_{\min}$, which determines the minimum amplitude of the ROS responses. That is, $g_{\min} = 1$ corresponds to no modulation whatsoever, so that sequence identity has no effect on ROS activity, whereas $g_{\min} = 0$ corresponds to the maximum modulation possible, in which case a neuron can be completely suppressed in its least preferred sequence. In general, a lower $g_{\min}$ value should give rise to stronger differentiation between sequences, higher signal-to-noise ratio, and thus higher accuracy (Salinas, 2004b). This is indeed what happens, as shown in Figure 4, *c* and *d*. The thin black lines in these plots indicate the standard condition in which $g_{\min} = 0.4$. With $g_{\min} = 0$ (orange lines), the stronger modulation indeed decreases the probability of coding the wrong movement, although the decrease is relatively modest, particularly in the case of P_M . In contrast, when $g_{\min} = 0.85$ (purple lines), the maximum suppression is only 15%. This small amount of modulation produces rather drastic increases in both P_m and P_M , which are comparable with the increases seen from 6 to 18 stored sequences in Figure 4, *a* and *b*. Therefore, to obtain high motor accuracy, the minimum gain $g_{\min}$ should be considerably below 1, but it does not have to be very close to zero.

Finally, consider the effect of synaptic corruption on the performance of the model, as measured by P_m and P_M . In this case, each model network is simulated as before except that, after training, each synaptic weight w_{ij} is randomly set to zero with a probability P_w . In other words, $P_w = 0.25$, for instance, means that $\sim 25\%$ of the synaptic connections are randomly deleted. The behavior of the model in the presence of such synaptic corruption is shown in Figure 4, *e* and *f*, where, again, the thin black lines correspond to the standard condition in which $P_w = 0$. Randomly deleting 5% of the network connections (orange lines) increases both P_m and P_M by small amounts, and the manipulation has a negligible effect on large networks. The increase is much more notorious when 25% of the connections are randomly set to zero (purple lines), but the result is not catastrophic, in the sense that the probability of error still goes down exponentially with the number of ROS neurons in the network. Therefore, the model is not overly sensitive to corruption of the synaptic connections.

Mastering an individual motor sequence

Experimental studies show that monkeys are capable of working with a relatively large repertoire of motor sequences (Shima and Tanji, 2000; Lu et al., 2002; Averbeck et al., 2006). In a typical motor sequence paradigm, at the start of each block of trials, the monkey must execute a new sequence whose identity has to be discovered through trial and error. After some number of exploratory trials, the sequence is discovered and the monkey performs

it a few more times from memory alone. Monkeys become more efficient in this process after a period of extensive training during which they become familiar with the available repertoire of sequences. Therefore, performance in these tasks depends in part on the learned repertoire, but in addition, every time a new block starts, there are a few initial trials during which the current sequence is recalled and quickly mastered (Procyk et al., 2000; Lee and Quessy, 2003).

The network model simulates the generation of a number of learned sequences when these need to be reproduced from memory alone, but what happens when one sequence (the current one) needs to be performed more accurately or is practiced a few times, as in the monkey experiments? In particular, if the overall motor accuracy for the set of learned sequences is not very high, how much does the network need to change to improve the performance of one specific sequence, and does this have an impact on the other sequences?

These questions are addressed by including in the simulations a set of coefficients ϕ_q that tell the network how important sequence q is (see Materials and Methods). Two situations are contrasted: the standard case in which all sequences have equal importance, and a second case in which sequence 1 is singled out and is given an importance coefficient that is different from all the others (it does not matter which sequence is singled out; sequence 1 is chosen for convenience). The importance coefficients are such that $\sum_q \phi_q = 1$, so notice the following: (1) $\phi_1 = 1$ means that sequence 1 is the only one that matters, all others are ignored, (2) $\phi_1 = 0$ means that sequence 1 is ignored, and (3) $\phi_1 = 1/N_Q$, where N_Q is the number of available sequences, means that all sequences have the same importance. Therefore, when $\phi_1 > 1/N_Q$, sequence 1 is more important than any of the others and, when $\phi_1 < 1/N_Q$, sequence 1 is less important.

Figure 5 illustrates what happens when ϕ_1 is varied. In this example, the network included 252 ROS and 6 motor neurons and was trained to reproduce 18 sequences of three movements. Figure 5*a* shows the mean responses of the six motor neurons in the network, as functions of time, for 6 of the 18 sequences. These are the responses in the standard condition in which all sequences are equally important. The mean probability of encoding an incorrect movement in this case was $P_m = 0.15$, but the probability varied slightly across sequences: the three sequences at the top of Figure 5*a* had the highest values, whereas the three sequences at the bottom had the lowest. Figure 5*b* shows the motor responses of the same network under identical conditions except that sequence 1, AAB, had a higher importance. Notice that sequence AAB is reproduced much more accurately in Figure 5*b* than in Figure 5*a*: only the neurons that are supposed to be active fire intensely at each time point. In contrast, the rest of the sequences are reproduced only slightly less accurately than in Figure 5*a*. This means that the network can generate a large improvement in one specific sequence with a relatively small cost in accuracy in each of the other sequences. A crucial question, however, is whether this improvement requires drastic changes in the synaptic connections between ROS and motor units.

It turns out that the answer is no, as illustrated in Figure 5*c*. The synaptic weights w_{ij}^* that result when sequence 1 is favored have very similar values as the synaptic weights w_{ij} that result when sequence 1 has the same importance as all the others; the correlation coefficient between them is $\rho = 0.985$. The similarity between $\mathbf{w}^*$ and $\mathbf{w}$ results not because sequence 1 is special but rather because of the properties of the network: each motor response is driven by many ROS neurons, so increasing the accuracy of one particular sequence requires many small changes in

Figure 5. Improvement in the representation of one selected sequence among many. A network of 252 ROS and 6 motor cells was trained to generate 18 sequences of three movements under two conditions: with equal importance for all sequences ($\phi_1 = 1/18$), or with a different importance for the first sequence and equal importance for the rest ($\phi_1 \neq 1/18$). *a*, Firing rates of the six motor neurons as functions of time obtained with equal importance. The three sequences shown at the top had the highest E_{RMS} values; the three sequences at the bottom had the lowest values. *b*, Firing rates of the six motor neurons as functions of time when the first sequence, AAB, had a higher importance ($\phi_1 = 1/5$). *c*, Synaptic weights from the simulation in *b* versus the synaptic weights from the simulation in *a*. Each dot indicates one of the 252×6 connections in the network. The correlation coefficient is $\rho = 0.985$. *d*, Motor accuracy and synaptic correlation as functions of ϕ_1 . The red curve indicates the probability of motor error P_m for the first sequence (AAB). The dark blue curve indicates P_m averaged over all sequences other than the first one. The light blue curve indicates the correlation coefficient ρ between synaptic weight values. The vertical dashed line on the left marks the point of equal importance; the dashed line on the right marks the value of ϕ_1 used in *b* and *c*.

synaptic weight rather than a few large changes. The impact of ϕ_1 on the network can be seen more clearly in Figure 5*d*, which plots three quantities: the probability of encoding a motor error during sequence 1 (red trace), the probability of encoding a motor error during any of the other sequences (dark blue trace), and the correlation coefficient between $\mathbf{w}^*$ and $\mathbf{w}$, all as functions of ϕ_1 . Interestingly, the shapes of these curves are ideal for changing the accuracy of one particular sequence without affecting the others or the synaptic weights very much. Around the point of equal importance ($\phi_1 = 1/18$, indicated by the dashed vertical line on the left), the red curve is very steep, which means that small changes in ϕ_1 will translate into large changes in the accuracy with which sequence 1 is executed. In contrast, around the point of equal importance, the light and dark blue curves are very flat, which means that small changes in ϕ_1 around this point will have a minimal impact on the other sequences and on the synaptic weights. As one would expect, increasing ϕ_1 eventually does alter the weights drastically and gives rise to large motor errors during all other sequences, but this happens for values larger than ~ 0.6 , which are extremely large.

Figure 6. Effects of inactivating and stimulating a subset of similar ROS neurons. The firing rates of three motor neurons in the network are plotted as functions of time. The network (same as in Figs. 1, 2) was trained to reproduce the six sequences indicated in each column. **a**, Standard condition. Same data and format as in Figure 2b. **b**, Inactivation condition. Two-thirds of the ROS neurons that fire during the second preparatory period were inactivated by setting their firing rates to 40% of their normal values. Note the decreased motor responses during this period (arrow). **c**, Stimulation condition. The same ROS neurons were then stimulated by adding 30 spikes/s to their normal firing rate values. This shifted the motor firing rates by 10.9 (magenta), 7.3 (light blue), and -0.2 (dark blue) spikes/s.

These effects, the steepness of the curve for sequence 1 and the flatness of the other two curves, become stronger in larger networks (results not shown). Thus, the organization of these networks is such that a large improvement in the execution of one particular sequence can be achieved through very small changes in synaptic connectivity and with minimal consequences for the execution of other sequences.

Experimental predictions

With the model, it is possible to simulate inactivation and stimulation experiments that would probe the connectivity of the networks involved in sequence generation. Two such manipulations were performed on the same model network used in Figures 1 and 2.

Inactivation of a homogeneous subset of ROS neurons was simulated first. The targeted ROS neurons were two-thirds of those that fired during the preparatory period preceding the second movement. These cells were inactivated by multiplying their firing rates by 0.4, whereas the rest of the ROS neurons fired at

their usual rates (as in Fig. 1). Inactivation was applied throughout the whole simulation time. Not surprisingly, the result was a reduction in firing rate in most motor neurons during the second preparatory period. This can be seen by comparing the three example motor responses in Figure 6a, the standard condition, with those in Figure 6b, the inactivation condition. The motor responses during the second preparatory period were suppressed by $\sim 50\%$, although the exact amount varied across cells. The key, however, is that the effect was temporally specific: inactivation of the ROS neurons that signal the second preparatory period had an impact on the motor responses exclusively during that period.

The second manipulation is more interesting. The same subset of ROS neurons was stimulated by adding 30 spikes/s to their usual firing rates at all times. In this way, these model neurons became active at the “wrong” times, at which they would normally have been silent. As can be seen in Figure 6c, this typically shifted the firing rate curves of the downstream motor neurons upward. The size and direction of the shift varied across motor cells, but the key is that the effect was not temporally specific: for

a given cell, the shift was the same at all times and for all sequences.

Similar shifts were obtained when the same number of ROS cells were stimulated but these were chosen randomly from the population. That is, the effect of such stimulation on the motor responses is qualitatively the same whether the stimulated ROS cells have the same rank-order selectivity or not. This is important because it means that an actual stimulation experiment could work even if the local stimulated ROS population contains a variety of temporal selectivities.

The observed shifts happen, of course, because the synaptic connections between ROS and motor neurons have fixed strengths, and thus the applied stimulation is passed on to the motor responses simply scaled by some amount. What changes across sequences, and which allows the motor neurons to fire at different times for different sequences, are the relative amplitudes of the ROS responses. Therefore, the stimulation and inactivation experiments potentially provide a separate method for verifying the crucial feature of the model, the nonlinear combination of rank-order and sequence selectivities of the ROS neurons.

Discussion

The neural network model just described can generate a variety of motor sequences. It achieves this by using a type of neuronal response that has two principal characteristics: (1) for each cell, strong activation occurs over a specific period of time that is fixed across different sequences, and (2) the overall response amplitude, or gain, of each cell varies across sequences. The first property is solidly established experimentally; the second is true at least in some of the published examples but has not been analyzed rigorously. According to the model, however, this second property is key, so the responses of most ROS neurons should display a perhaps modest but clearly significant sensitivity to sequence identity.

The network is also robust to synaptic corruption and behaves in a way that is ideally suited for mastering one particular sequence quickly. Thus, when a current sequence of movements becomes important, a large increase in accuracy can be achieved through small changes in synaptic connectivity, which could conceivably take place within the time span of a few practice trials (Procyk et al., 2000; Lee and Quessy, 2003), without wrecking the performance of other stored sequences.

Coordinate transformations in space and time

Functionally, the modulation by sequence identity exploited here is akin to the modulation by gaze angle that has been documented in parietal cortex (Andersen et al., 1985, 1990, 1993; Brothie et al., 1995) and that is thought to be crucial for representing the locations of objects with respect to different body parts or with respect to the world (Zipser and Andersen, 1988; Salinas and Abbott, 1995; Andersen et al., 1997; Pouget and Sejnowski, 1997; Snyder et al., 1998; Buneo et al., 2002). In parietal cortex, the modulation is by a proprioceptive signal, the responses are triggered by visual stimuli, and the key variable is spatial location. The situation seems completely different in the SMA and pre-SMA, in which the modulation is by sequence identity, the ROS activity is driven by an internal representation of elapsed time (or number of elapsed events), and the main variable is time. Computationally, however, the underlying mechanisms are very similar: the gain-modulated neurons have a fixed preferred value along the relevant axis (space or time), whereas the downstream neurons have responses that can switch positions along the relevant axis (space or time). This transformation is a consequence of

what is described mathematically as a basis–function representation (Poggio, 1990; Pouget and Sejnowski, 1997). Such representations have been thoroughly studied in the spatial domain (Salinas and Thier, 2000; Xing and Andersen, 2000; Deneve et al., 2001; Deneve and Pouget, 2003; Salinas, 2004a,b) but may also be fundamental in the time domain (Botvinick and Watanabe, 2007) (see also Wainwright et al., 2005).

What have we learned?

Some of the original experimental reports documenting ROS responses suggested that they could serve as building blocks for assembling motor sequences (Mushiake and Strick, 1995; Shima and Tanji, 2000; Tanji, 2001). These simulations, however, are the first quantitative demonstration of this idea. In addition, they clarify three features of the data that seemed rather mysterious.

First, why are there at least some instances in which the ROS activity changes with the sequence? Without sequence-specific modulation, ROS neurons can only generate a fixed output, which means a single sequence; it is the modulation that makes them powerful and allows them to act as a temporal basis set, i.e., as building blocks for constructing multiple arbitrary sequences of movements downstream.

Second, some units show a strong dependence on sequence identity, whereas others seem to be entirely insensitive. Why are there so many types of ROS neurons? First, a wide diversity in modulation factors does not imply that there are fundamentally different cell classes; they could all participate in the same function. Different neurons may be simply modulated to different extents, or perhaps this results because the full modulation range of a cell can only be determined with a much larger repertoire of sequences. Second, diversity of modulation factors is in no way a problem for the performance of the circuit. What is important is that neurons should display all combinations of sequence and rank preferences.

Third, very few ROS neurons show two or more peaks of activity; why? In the experiments, some neurons were active in more than one, isolated epoch during a sequence (Shima and Tanji, 2000, their Fig. 6). A neuron, for instance, could be strongly active both during the first and the third preparatory periods. In the model, with no additional constraints, this makes virtually no difference; in principle, the same repertoire of downstream responses can be generated with single or multiple activation periods per neuron, as long as all the relevant combinations of preferred time segments (single or multiple) and preferred sequences are present. However, having ROS neurons that respond during single time intervals in the sequences drastically simplifies the calculation of the synaptic weights between them and the motor neurons (see Appendix). Therefore, the observed representation, in which most ROS units are active only during one discrete part of a sequence, is likely to be particularly convenient for learning.

In addition, two specific predictions arise from the general framework of neural basis functions (Salinas and Abbott, 1995, 1997; Pouget and Sejnowski, 1997; Ben Hamed et al., 2003; Salinas, 2004b) applied to this case. (1) The effect of sequence identity on the temporal profiles must be nonlinear, as in Equation 1, in which it is multiplicative. This means it cannot be simply additive, for example (as in Eq. 5). If it is, the model does not work at all: for instance, for the network in Figure 1, P_M increases from 0 to 0.55. (2) Preferred sequences and preferred activation time periods (ranks) should be distributed independently. The appropriate analyses for verifying these predictions remain to be done.

Sequences of sensory and motor events

There are three types of neurophysiological paradigms involving sequences: (1) a subject must produce a complete sequence of motor actions from memory, (2) a subject must identify or discriminate sequences of stimuli presented in a given order (e.g., is your phone number 314-159-2653?), or (3) a subject must produce a sequence of motor actions such that each one is either accompanied or guided by a sensory stimulus. Case 1 corresponds to the purely motor tasks investigated here, and the proposal is that they depend on ROS responses modulated by sequence identity.

Case 2 is a perceptual task. Botvinick and Watanabe (2007) recently showed that such a task can be performed on the basis of populations of neurons that are selective for individual sensory stimuli (shapes or locations) and are modulated by rank; that is, by the number of items that have been presented so far. Neurons with such combined selectivities have indeed been found in the prefrontal cortex and basal ganglia (Kermadi and Joseph, 1995; Ninokura et al., 2003, 2004; Inoue and Mikami, 2006). Thus, although the selectivities involved are very different, the circuits proposed here and by Botvinick and Watanabe rely on similar mechanisms for integrating information and generating a transformed representation downstream.

Case 3, conversely, is a mixed situation in which a sequence of motor actions is performed but not from memory alone. Neurophysiological studies using such tasks (Clower and Alexander, 1998; Lu et al., 2002; Ninokura et al., 2003, 2004; Averbach et al., 2006) have reported neurons with either straightforward motor properties or with combinations of rank-order (temporal) and sequence selectivities, as in the memory-based tasks of Tanji and colleagues (Ninokura et al., 2004; Averbach et al., 2006). Interestingly, however, in some cases, an intermediate motor representation seems to develop, whereby neurons prefer a specific motor action but only when it is triggered by a particular stimulus (Clower and Alexander, 1998; Lu et al. 2002). For example, a neuron could be active only when the display shows two dots aligned vertically and the target is the one on top. Such activation is not rank selective but could again result from a display-specific preference (for instance, to two vertically aligned dots) combined with a sequence-specific preference (Salinas, 2004a,b). Task-dependent sensory responses of similar complexity have been reported previously (Lauwereyns et al., 2001; Koida and Komatsu, 2007), so this is a reasonable possibility.

Mechanisms underlying sequence generation

In general, the generation of serial behavior encompasses many tasks and is likely to involve multiple neural mechanisms and representations. In the case analyzed here, modulated ROS responses arise when arbitrary sequences of discrete, unrelated movements are executed. In contrast, when monkeys are trained to draw geometrical shapes, prefrontal–cortex responses encoding each segment of a drawing are activated in parallel rather than during discrete periods (Averbach et al., 2002, 2003). Models based on this type of parallel activity account for a variety of serial-order effects observed in skilled behaviors, such as speech production and cursive handwriting (Bullock, 2004; Rhodes et al., 2004).

Human psychophysical studies of sequence generation often use the immediate serial recall task, in which a list of items is presented and a subject must recall all items in the original order. Error patterns in this paradigm show a number of regularities having to do with list length, item repetition rates, and item discriminability, and modeling studies suggest that these effects de-

pend strongly on how the list is represented and dynamically updated in working memory (Rhodes et al., 2004; Botvinick and Plaut, 2006). However, because such models aim to explain a wide range of behavioral experiments, they are more abstract than the present one and revolve around more general issues, for instance, whether a recurrent network architecture can, in principle, account for all the data.

Unlike these models, the current framework is based on a specific neuronal representation and analyzes one experimentally identified step in the computational machinery required for producing a particular type of motor sequence. The next question in this regard is how the ROS cells are constructed in the first place. The underlying process is likely to require one or more mechanisms for time integration (Dehaene et al., 1987; Drew and Abbott, 2003; Mauk and Buonomano, 2004; Jin et al., 2007; Karmarkar and Buonomano, 2007). This will be an important challenge for future work.

Appendix

This appendix describes a simplified version of the model with which it can be shown analytically (1) that the minimum number of ROS neurons needed to store N_Q distinct sequences is $N_{ROS} = N_Q \times N_S$, where N_S is the number of steps or time periods per sequence, and (2) that having ROS neurons that are activated during a single time period simplifies the learning process.

For simplicity, assume (1) that $\alpha = 0$, so there is no response variability, (2) that the background firing rate is zero, (3) that all sequences have the same importance, and (4) that there are N_S segments or steps per sequence and each neuron is either active or inactive during each one of those steps; in other words, each movement or preparatory period is considered as one step, and a neuron, whether ROS or motor, fires at a constant rate throughout a step. Then, assuming multiplicative modulation by sequence identity, the firing rate of ROS neuron j during step s of sequence q is

$$r_{jqs} = g_{jq} f_{js}, \quad (14)$$

where j goes from 1 to N_{ROS} , q goes from 1 to N_Q , and s goes from 1 to N_S .

In the examples discussed in Materials and Methods, there are $N_Q = 6$ sequences, each composed of three movement periods, three preparatory periods, and a trailing blank period at the end, so the number of steps is $N_S = 7$ for those cases.

Now let us find an exact solution to the problem, that is, a set of ROS responses and a set of synaptic weights such that the driven motor responses are exactly equal to the desired ones. This means that the following equality should be satisfied:

$$\sum_j w_{kj} r_{jqs} = M_{kqs} \quad (15)$$

where M_{kqs} is the desired response of motor neuron k during step s of sequence q , and the expression on the left side is the corresponding driven response. Note that, in this equation, the indices q and s appear together on both sides. Therefore, they can be replaced by a single index i that runs from 1 to $N_Q \times N_S$. That is,

$$\sum_j w_{kj} r_{ji} = M_{ki} \quad (16)$$

is equivalent to Equation 15. This, however, is a standard matrix equation, $\mathbf{w}\mathbf{r} = \mathbf{M}$. Therefore, an exact solution for $\mathbf{w}$ can be found for any $\mathbf{M}$ as long as the number of independent rows of $\mathbf{r}$

equals or exceeds its number of columns (Barnett, 1990), that is, as long as $N_{\text{ROS}} \geq N_Q \times N_S$. This proves claim 1.

A couple of examples may provide some additional insight. If $\mathbf{r}$ is square ($N_{\text{ROS}} = N_Q \times N_S$), and its rows are linearly independent (i.e., the ROS cells respond differently from each other beyond additive and multiplicative constants), then the inverse of $\mathbf{r}$ exists and the synaptic weights that solve the problem are

$$\mathbf{w} = \mathbf{M}\mathbf{r}^{-1} \quad (17)$$

and they are unique. In contrast, if $N_{\text{ROS}} < N_Q \times N_S$, then weights can be found that approximate Equation 16 in the least-squares sense, but there will be some sequences that the network will not be able to generate accurately. Finally, if $N_{\text{ROS}} > N_Q \times N_S$, then there are many possible solutions, one of which is

$$\mathbf{w} = \mathbf{L}\mathbf{C}^{-1} \quad (18)$$

with

$$C_{jk} \equiv \sum_{q,s} r_{jq} r_{ks} \quad (19)$$

$$L_{jk} \equiv \sum_{q,s} M_{jq} r_{ks} \quad (20)$$

which is the same type of solution as in Equations 10–12.

To prove claim 2, first substitute Equation 14 into Equation 15, so that

$$\sum_j w_{kj} g_{jq} f_{js} = M_{kq} \quad (21)$$

and consider what happens when each ROS unit is active only during one step (the same in all sequences). Suppose that the temporal profile of ROS unit j is

$$f_{js} = \begin{cases} 1 & \text{if unit } j \text{ fires during step } s \\ 0 & \text{otherwise} \end{cases} \quad (22)$$

In that case, only those ROS neurons that fire during step s contribute to the motor response M_{kq} . As a consequence, Equation 21 can be rewritten as follows:

$$\sum_{i=1}^{N_{\text{ROS}}(s)} w_{ki}^s g_{iq}^s = M_{kq}^s \quad (23)$$

where $N_{\text{ROS}}(s)$ is the number of ROS neurons that are active during step s , the index k now runs from 1 to the number of motor neurons that fire during step s , and the superscript s means “for neurons that fire during step s .” In effect, by doing this, Equation 21 is broken up into N_S independent matrix equations, one for each step. The above expression is equivalent to

$$\mathbf{w}^s \mathbf{g}^s = \mathbf{M}^s, \quad (24)$$

which shows that the problem is now local in time: only those ROS neurons that fire during step s drive the motor responses that occur during that step and contribute to the matrix $\mathbf{w}^s$. Thus, in practice, finding all the network connections now requires the solution to N_S small systems of equations (Eq. 24), which are decoupled. In contrast, without the temporal constraint imposed by Equation 22, each ROS neuron can fire during several steps, and as a consequence, a much larger, fully coupled system of equations must be solved simultaneously. Here is another way to think about the distinction: with Equation 24, the connections

from presynaptic neuron i are updated only when unit i is active, whereas in general, the connections from unit i must be updated at every step, even if unit i is silent. This is what happens in Equations 18–20 as well as in Equation 17.

This analysis does not reveal which biological synaptic modification mechanisms actually work to establish the necessary connections in reality (Davison and Frégnac, 2006), but whatever they are, the problem is vastly simplified if the correct strength of a synapse depends only on the presynaptic and postsynaptic activities evaluated at the same time rather than at different times.

References

- Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. *Science* 230:456–458.
- 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* 10:1176–1196.
- Andersen RA, Snyder LH, Li CS, Stricanne B (1993) Coordinate transformations in the representation of spatial information. *Curr Opin Neurobiol* 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.
- Ashe J, Lungu OV, Basford AT, Lu X (2006) Cortical control of motor sequences. *Curr Opin Neurobiol* 16:213–221.
- Averbeck BB, Chafee MV, Crowe DA, Georgopoulos AP (2002) Parallel processing of serial movements in prefrontal cortex. *Proc Natl Acad Sci USA* 99:13172–13177.
- Averbeck BB, Crowe DA, Chafee MV, Georgopoulos AP (2003) Neural activity in prefrontal cortex during copying geometrical shapes. II. Decoding shape segments from neural ensembles. *Exp Brain Res* 150:142–153.
- Averbeck BB, Sohn JW, Lee D (2006) Activity in prefrontal cortex during dynamic selection of action sequences. *Nat Neurosci* 9:276–282.
- Barnett S (1990) *Matrices: methods and applications*. Oxford: Oxford UP.
- 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.
- Berdyeva TK, Olson CR (2009) Monkey supplementary eye field neurons signal the ordinal position of both actions and objects. *J Neurosci* 29:591–599.
- Botvinick M, Watanabe T (2007) From numerosity to ordinal rank: a gain-field model of serial order representation in cortical working memory. *J Neurosci* 27:8636–8642.
- Botvinick MM, Plaut DC (2006) Short-term memory for serial order: a recurrent neural network model. *Psychol Rev* 113:201–233.
- 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.
- Bullock D (2004) Adaptive neural models of queuing and timing in fluent action. *Trends Cogn Sci* 8:426–433.
- Buneo CA, Jarvis MR, Batista AP, Andersen RA (2002) Direct visuomotor transformations for reaching. *Nature* 416:632–636.
- Clower WT, Alexander GE (1998) Movement sequence-related activity reflecting numerical order of components in supplementary and presupplementary motor areas. *J Neurophysiol* 80:1562–1566.
- Davison AP, Frégnac Y (2006) Learning cross-modal spatial transformations through spike timing-dependent plasticity. *J Neurosci* 26:5604–5615.
- Dehaene S, Changeux JP, Nadal JP (1987) Neural networks that learn temporal sequences by selection. *Proc Natl Acad Sci USA* 84:2727–2731.
- Deneve S, Pouget A (2003) Basis functions for object-centered representations. *Neuron* 37:347–359.
- Deneve S, Latham PE, Pouget A (2001) Efficient computation and cue integration with noisy population codes. *Nat Neurosci* 4:826–831.
- Drew PJ, Abbott LF (2003) Model of song selectivity and sequence generation in area HVC of the songbird. *J Neurophysiol* 89:2697–2706.
- Gerloff C, Corwell B, Chen R, Hallett M, Cohen LG (1997) Stimulation over the human supplementary motor area interferes with the organization of future elements in complex motor sequences. *Brain* 120:1587–1602.
- Haykin S (1999) *Neural networks: a comprehensive foundation*. Upper Saddle River, NJ: Prentice Hall.

- Inoue M, Mikami A (2006) Prefrontal activity during serial probe reproduction task: encoding, mnemonic and retrieval processes. *J Neurophysiol* 95:1008–1041.
- Isoda M, Tanji J (2003) Contrasting neuronal activity in the supplementary and frontal eye fields during temporal organization of multiple saccades. *J Neurophysiol* 90:3054–3065.
- Isoda M, Tanji J (2004) Participation of the primate presupplementary motor area in sequencing multiple saccades. *J Neurophysiol* 92:653–659.
- Jin DZ, Ramazanoglu FM, Seung HS (2007) Intrinsic bursting enhances the robustness of a neural network model of sequence generation by avian brain area HVC. *J Comput Neurosci* 23:283–299.
- Karmarkar UR, Buonomano DV (2007) Timing in the absence of clocks: encoding time in neural network states. *Neuron* 53:427–438.
- Kermadi I, Joseph JP (1995) Activity in the caudate nucleus of monkey during spatial sequencing. *J Neurophysiol* 74:911–933.
- Koida K, Komatsu H (2007) Effects of task demands on the responses of color-selective neurons in the inferior temporal cortex. *Nat Neurosci* 10:108–116.
- 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.
- Lee D, Quessy S (2003) Activity in the supplementary motor area related to learning and performance during a sequential visuomotor task. *J Neurophysiol* 89:1039–1056.
- Lu X, Matsuzawa M, Hikosaka O (2002) A neural correlate of oculomotor sequences in supplementary eye field. *Neuron* 34:317–325.
- Mauk MD, Buonomano DV (2004) The neural basis of temporal processing. *Annu Rev Neurosci* 27:307–340.
- Mushiake H, Strick PL (1995) Pallidal neuron activity during sequential arm movements. *J Neurophysiol* 74:2754–2758.
- Nachev P, Kennard C, Husain M (2008) Functional role of the supplementary and presupplementary motor areas. *Nat Rev Neurosci* 9:856–869.
- Ninokura Y, Mushiake H, Tanji J (2003) Representation of the temporal order of visual objects in the primate lateral prefrontal cortex. *J Neurophysiol* 89:2868–2873.
- Ninokura Y, Mushiake H, Tanji J (2004) Integration of temporal order and object information in the monkey lateral prefrontal cortex. *J Neurophysiol* 91:555–560.
- Peña JL, Konishi M (2001) Auditory spatial receptive fields created by multiplication. *Science* 292:249–252.
- Poggio T (1990) A theory of how the brain might work. *Cold Spring Harb 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.
- Procyk E, Tanaka YL, Joseph JP (2000) Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nat Neurosci* 3:502–508.
- Rhodes BJ, Bullock D, Verwey WB, Averbeck BB, Page MP (2004) Learning and production of movement sequences: behavioral, neurophysiological, and modeling perspectives. *Hum Mov Sci* 23:699–746.
- Salinas E (2004a) Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. *J Neurosci* 24:1113–1118.
- Salinas E (2004b) Context-dependent selection of visuomotor maps. *BMC Neurosci* 5:47.
- 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, Thier P (2000) Gain modulation: a major computational principle of the central nervous system. *Neuron* 27:15–21.
- Shima K, Tanji J (1998) Both supplementary and presupplementary motor areas are crucial for the temporal organization of multiple movements. *J Neurophysiol* 80:3247–3260.
- Shima K, Tanji J (2000) Neuronal activity in the supplementary and presupplementary motor areas for temporal organization of multiple movements. *J Neurophysiol* 84:2148–2160.
- 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.
- Sohn JW, Lee D (2007) Order-dependent modulation of directional signals in the supplementary and presupplementary motor areas. *J Neurosci* 27:13655–13666.
- Tanji J (2001) Sequential organization of multiple movements: involvement of cortical motor areas. *Annu Rev Neurosci* 24:631–651.
- Tanji J, Shima K (1994) Role for supplementary motor area cells in planning several movements ahead. *Nature* 371:413–416.
- Wainscott SK, Donchin O, Shadmehr R (2005) Internal models and contextual cues: encoding serial order and direction of movement. *J Neurophysiol* 93:786–800.
- Xing J, Andersen RA (2000) Models of the posterior parietal cortex which perform multimodal integration and represent space in several coordinate frames. *J Cogn Neurosci* 12:601–614.
- 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.