Context-Dependent Modulation of Functional Connectivity: Secondary Somatosensory Cortex to Prefrontal Cortex Connections in Two-Stimulus-Interval Discrimination Tasks

Stephanie S. Chow,² Ranulfo Romo,¹ and Carlos D. Brody^{2,3}

¹Instituto de Fisiología Celular-Neurociencias, Universidad Nacional Autónoma de México, 04510 México City, México, and ²Howard Hughes Medical Institute, Princeton Neuroscience Institute and ³Department of Molecular Biology, Princeton University, Princeton, New Jersey 08544

In a complex world, a sensory cue may prompt different actions in different contexts. A laboratory example of context-dependent sensory processing is the two-stimulus-interval discrimination task. In each trial, a first stimulus (f_1) must be stored in short-term memory and later compared with a second stimulus (f_2), for the animal to come to a binary decision. Prefrontal cortex (PFC) neurons need to interpret the f_1 information in one way (perhaps with a positive weight) and the f_2 information in an opposite way (perhaps with a negative weight), although they come from the very same secondary somatosensory cortex (S2) neurons; therefore, a functional sign inversion is required. This task thus provides a clear example of context-dependent processing.

Here we develop a biologically plausible model of a context-dependent signal transformation of the stimulus encoding from S2 to PFC. To ground our model in experimental neurophysiology, we use neurophysiological data recorded by R. Romo's laboratory from both cortical area S2 and PFC in monkeys performing the task. Our main goal is to use experimentally observed context-dependent modulations of firing rates in cortical area S2 as the basis for a model that achieves a context-dependent inversion of the sign of S2 to PFC connections. This is done without requiring any changes in connectivity (Salinas, 2004b). We (1) characterize the experimentally observed context-dependent firing rate modulation in area S2, (2) construct a model that results in the sign transformation, and (3) characterize the robustness and consequent biological plausibility of the model.

Introduction

In a constantly changing world, processing of signals from the environment must be flexible: the same sounds that in one context (e.g., our own home) may lead us to pick up a telephone may, in a different context (e.g., someone else's home) lead us to merely look expectantly at our host (Miller and Cohen, 2001; Salinas, 2004a). A laboratory example of context-dependent sensory processing can be found in commonly used two-stimulus-interval discrimination tasks. In each trial, a first stimulus must be stored in short-term memory, whereas a second stimulus must instead be compared with the memory of the first and used to come to a binary decision. How are these two stimuli treated differently?

Our laboratory recently proposed a network model of processing in the prefrontal cortex (PFC) that addresses short-term memory and decision-making in two-stimulus-interval discrimination tasks (Machens et al., 2005), but the model requires con-

nections from the secondary somatosensory cortex (S2) to PFC to switch sign between the first vibrotactile stimulus period (f_1) and the second (f_2). This sign inversion is part of how the two stimuli are treated differently and is crucial to the proposed model, yet how it is achieved in biology is unclear. Here we propose a biologically plausible alternative that builds on the insights of Salinas' models of fast switching between motor actions (Salinas, 2004a,b).

Salinas' central idea was that, when a population of sensory neurons responds to stimuli in a manner that depends not only on the current stimulus but also on the current context, then simple linear weighted sums of neuronal activities can transform the existing context dependence of the sensory neurons into a desired function of context.

In the present case, the desired function is a sign inversion in the encoding of stimulus frequency between the first and second stimulus. We first analyze the patterns of experimentally observed firing rates in S2 to ascertain the degree of context-dependent modulation and find context dependence in a small but significant fraction of neurons. Context dependence allows us to build a simple yet plausible network model whose output grows linearly with frequency during f_1 yet decreases linearly with frequency during f_2 . Our proposed model would support a sufficiently high signal-to-noise ratio to be consistent with the known performance accuracy of the monkeys, and it is robust to perturbations in its connections. We are motivated by (1) the fact that the two-stimulus-interval task provides a clear example of

Received Oct. 5, 2008; revised Jan. 21, 2009; accepted Feb. 17, 2009.

R.R. was supported by an International Scholars Award from the Howard Hughes Medical Institute and by grants from the Dirección General del Personal Académico de la Universidad Nacional Autónoma de México and the Consejo Nacional de Ciencia y Tecnología. S.S.C. was supported in part by a fellowship from the Swartz Foundation. S.S.C. and C.D.B. were supported in part by National Institute of Mental Health Grant R01-MH067991. We thank A. Hernández, L. Lemus, and A. Zainos for technical assistance.

Correspondence should be addressed to Stephanie S. Chow at her present address: 18 Post Horn Place, Waterloo, Ontario N2L 5E9, Canada. E-mail: stephanieschow@gmail.com.

DOI:10.1523/JNEUROSCI.4856-08.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/297238-08\$15.00/0

Figure 1. The two-interval vibrotactile discrimination task. A mechanical probe is lowered until it indents the skin of one digit of the hand. The sequence of events is as follows: the probe delivers the first vibrotactile stimulus at the frequency f_1 for 500 ms; there is a brief delay, most often 3 s; the second stimulus is delivered at the frequency f_2 for 500 ms; finally, the monkey uses his free hand to press one of two buttons to indicate the result of the comparison.

context-dependent processing and is therefore an excellent substrate to address this general problem, and (2) the availability of neurophysiological data recorded by R. Romo's laboratory from both cortical areas S2 and PFC in monkeys performing the task, thus allowing a close grounding in experimental neurophysiology.

Materials and Methods

We use a dataset (Salinas et al., 2000; Romo et al., 2002) collected from monkeys performing a two-stimulus-interval vibrotactile discrimination task. Details of the experimental setup have been described previously (Hernández et al., 1997; Romo et al., 1998; Salinas et al., 2000). The two-stimulus-interval task depicted in Figure 1 requires both short-term memory and decision-making (Romo and Salinas, 2003) and was the basis for our model in which a context-dependent sign switch of sensory inputs to PFC was required (Machens et al., 2005).

During f_1 , many, although not all, neurons in PFC and S2 exhibit firing rates that are monotonic functions of the value of f_1 (Romo et al., 1999, 2002; Salinas et al., 2000; Brody et al., 2003). These neurons can be classified into two groups. Neurons that increase their firing rates with increases in f_1 are classified as “plus” (+) neurons, whereas those that decrease their firing rates with increasing stimulus frequency are classified as “minus” (−) neurons.

Most neurons in area S2 that are classified as plus during f_1 are also classified as plus during f_2 . Similarly, most neurons in area S2 that are classified as minus during f_1 are also classified as minus during f_2 . Remarkably, however, many neurons in PFC that are classified as plus during f_1 are classified as minus during f_2 ; conversely, many PFC neurons classified as minus during f_1 are classified as plus during f_2 (Machens et al., 2005). This suggests a switch, in the interval between f_1 and f_2 , in the sign of the functional connectivity from S2 to PFC. The focus of this study is on the question of how such a sign switch could be achieved and, in particular, whether differences in S2 responses during f_1 (the “ f_1 context”) and during f_2 (the “ f_2 context”) would be sufficient to account for the switch.

We use data from two monkeys (monkeys R13 and R14) during f_1 and f_2 to determine the effect of stimulus context on the stimulus responses. We concentrate on “task-dependent” neurons, those whose behavior changes after stimulus onset. Task-dependent neurons are those with significantly different firing rates after the onset of at least one of the two stimuli as determined by a t test comparing spike counts during the 500 ms before stimulus onset (whether f_1 or f_2) to spike counts during the 500 ms of stimulus presentation. All statistical tests are based on a significance level of 5%. We find 218 task-responsive neurons in monkey R13, and 704 in monkey R14 for a total of 922.

Of the 922 task-responsive S2 neurons from the Romo laboratory data, 893 (97%) have f_1 responses well fitted by a linear function over the vibrotactile stimulus range as indicated by a Q value > 0.01 in a χ^2 goodness-of-fit test (Press et al., 1992), whereas 850 (92%) have linear f_2 responses, and 834 (90%) have both linear f_1 and linear f_2 responses. Consequently, we proceed with the simplest linear model of the firing rate response (x) of an S2 neuron that allows F tests for context z and stimulus frequency f dependence, given by Equation 1:

$$x = \beta_0 + \beta_1 f + \beta_2 z + \beta_3 f z, \quad (1)$$

such that $z = 1$ means that the stimulus is presented during f_2 and $z = 0$ such that the stimulus is presented during f_1 .

The null hypothesis of no context dependence and the associated alternative hypotheses are specified by Equations 2 and 3. Of the 922 neurons, 111 of 922 (12.0%) are context dependent:

$$H_0: \beta_2 = \beta_3 = 0; \quad (2)$$

$$H_1: \text{Not } \beta_2 = \beta_3 = 0. \quad (3)$$

The null hypothesis of no f_1 dependence and its alternative hypotheses are specified by Equations 4 and 5:

$$H_0: \beta_1 = 0; \quad (4)$$

$$H_1: \text{Not } \beta_1 = 0. \quad (5)$$

The null hypothesis of no f_2 dependence and its alternative hypotheses are specified by Equations 6 and 7:

$$H_0: \beta_1 + \beta_3 = 0; \quad (6)$$

$$H_1: \text{Not } \beta_1 + \beta_3 = 0. \quad (7)$$

We consider a neuron to be frequency dependent if it is either or both f_1 and f_2 dependent. Of the 922 neurons, 361 (39.2%) are frequency dependent. Table 1 details the numbers of observed context- and frequency-dependent neurons by a monkey.

We use the S2 neuron data for inputs to our one-layer network firing rate model of Figure 5 that achieves the modulation of functional connectivity. If there are N S2 input neurons to the network, and it is trial t , the output of the network during f_1 , $y_t \in \mathbb{R}^+$, corresponds to the output of a neuron that is an input to PFC but is not necessarily in PFC. y_t is the weighted sum $\mathbf{W} \cdot \mathbf{x}_t$ of the inputs $\mathbf{x} = (x_{t1}, x_{t2}, \dots, x_{tN}) \in \mathbb{R}_N^+$ according to the synaptic weight vector $\mathbf{W} = (W_1, W_2, \dots, W_N) \in \mathbb{R}_N$. Similarly, the output of the network during f_2 , $\tilde{y}_t \in \mathbb{R}^+$, is the weighted sum $\mathbf{W} \cdot \tilde{\mathbf{x}}_t$ of the inputs $\tilde{\mathbf{x}}_t = (\tilde{x}_{t1}, \tilde{x}_{t2}, \dots, \tilde{x}_{tN}) \in \mathbb{R}_N^+$ according to the same synaptic weight vector $\mathbf{W}$ as during f_1 .

If T indicates the total number of trials, then the combination of weighted sums over all trials can be written in the matrix form of Equation 8 or more compactly in the form of Equation 9:

$$[W_1 \dots W_N] \begin{bmatrix} x_{11} & \tilde{x}_{11} & x_{T1} & \tilde{x}_{T1} \\ \dots & \dots & \dots & \dots \\ x_{1N} & \tilde{x}_{1N} & x_{TN} & \tilde{x}_{TN} \end{bmatrix}; \quad (8)$$

$$\mathbf{W} \mathbf{X} = \mathbf{Y}. \quad (9)$$

Our problem is to relate the inputs $\mathbf{X}$ to the desired outputs $\mathbf{Y}$ by finding an appropriate network, i.e., by finding appropriate weights $\mathbf{W}$.

If the rank of matrix $\mathbf{X}$ is less than or equal to N , the number of weights, then the number of unknowns is less than or equal to the number of equations and there exists at least one set of weights $\mathbf{W}$ that exactly maps the input $\mathbf{x}$ to the desired output $\mathbf{Y}$. If there is more than one solution, we choose the one that minimizes the sum of squared weights. If there is no exact solution, we find the solution that minimizes the sum of squared differences between the desired output and the actual output.

The intuition underlying our simple network is illustrated by the simple examples in Figure 2. Geometrically, the set of weights $\mathbf{W}$ that properly maps a given N neuron input vector $\mathbf{x}$ to its desired output y forms an N one-dimensional hyperplane. Each new $\mathbf{x}$ – y pair defines a new N one-dimensional hyperplane, and the intersection of all the hyperplanes, if it exists, is the set of all possible synaptic weights that correctly relates each input to its output. Thus, each additional $\mathbf{x}$ – y pair reduces the dimension of the intersection by one unless the new vector of responses $\mathbf{x}$ is a linear combination of some other input vectors, an unlikely situation when dealing with real data. Nonetheless, two very similar, and thus nearly collinear, input vectors $\mathbf{x}$ and $\mathbf{x}^*$ would result one or more very large weights as small differences in firing must be magnified to produce the

Table 1. Counts of task-dependent neurons by context dependence and frequency dependence in S2 from two monkeys, R13 and R14

		Not frequency dependent		Frequency dependent		Individual totals	
		R13	R14	R13	R14	R13	R14
Not context dependent	R13	124 (57%)		63 (29%)		187 (86%)	
	R14		385 (55%)		239 (34%)		624 (89%)
Context dependent	R13	17 (7.8%)		14 (6.4%)		31 (14%)	
	R14		35 (5.0%)		45 (6.4%)		80 (11%)
Individual totals	R13	141 (65%)		77 (35%)		218 (100%)	
	R14		420 (60%)		284 (40%)		704 (100%)

appropriate outputs. The resulting network would be very sensitive to noise in the neurons corresponding to large synaptic weights. In our network, more context dependence in S2 results in smaller weights and increased robustness to noise. One of our goals in this study is to quantify that robustness for the case of the S2→PFC data.

We calculate synaptic weights from the data in two ways: with mean firing rates and with per trial firing rates. In both cases, we add a bias term to the vector of inputs to account for non-task-dependent neurons. To measure network performance, we create a training set and a test set from the data, in which the training set is used to determine the synaptic weights and the test set is used to check how well the network generalizes to new data. For each neuron and for each stimulus, whether f1 or f2, we randomly assign two-thirds of the trials to be training trials, whereas the remaining one-third became test trials. We also pretend that our set of S2 neurons have all been simultaneously recorded from a single monkey so that we can combine multiple neurons into a single set of S2 inputs. Each resulting input data point is a vector of firing rates at a particular frequency, with one element per neuron. By repeatedly and independently sampling each neuron and stimulus over the training or test trials, we generate a training set and a test set.

Results

Based on both neurophysiological and neuroanatomical data (Courtney et al., 1998), it is thought that the working memory and decision-making processes required for the two-stimulus-interval vibrotactile discrimination task are subserved by frontal cortices, such as PFC. These frontal areas receive sensory input from S2. We show the following: (1) that significant context-dependent modulations are found in S2 during the discrimination task; (2) that, given appropriate connection strengths, these modulations are sufficient to change the functional connectivity between S2 to PFC and that this transformation is robust to variability in both neuronal firing rates and connection strengths; and, finally, (3) that this transformation leads to a signal-to-noise ratio high enough that it can be used to perform the task with an accuracy equal to or greater than that observed in experiments.

When the second stimulus (f2) is presented, neuronal activity in PFC evolves toward activity that is correlated with the binary decision the animal must make (Romo et al., 2002; Machens et al., 2005). In area S2, initial activity after the onset of stimulus f2 reflects the value of f2 but later evolves, as in PFC, to become dependent on f2–f1 rather than on f2 alone (Romo et al., 2002, 2004). This evolution toward (f2–f1) dependence occurs earlier in PFC (~150 ms after the onset of f2) (J. K. Jun, C. D. Brody, and R. Romo, unpublished results) than in S2 (181 ms after the onset of f2) (Romo et al., 2004). For this reason, it is thought that PFC

Figure 2. *A*, A simple two-dimensional example of synaptic weight determination. A vector x of S2 input firing rates during f1 and its corresponding desired output firing rate y are related by an unknown vector of synaptic weights W by the equation $W \cdot x = y$, thus defining a line. During f2, the new vector of inputs $\tilde{x}$ and the new desired f2 output $\tilde{y}$ are related by the same vector of synaptic weights W by the similar equation $W \cdot \tilde{x} = \tilde{y}$, thus defining a second line. The intersection W^* of the two lines gives the unique pair of weights that correctly relates the inputs to their desired outputs. *B*, An example in which extremely large weights would result when input firing rate vectors x and $\tilde{x}$ are nearly collinear, but their desired outputs y and $\tilde{y}$ are quite different.

(and possibly other frontal areas) are responsible for computing the monkey's decision and that, during the initial portion of stimulus f2, the role of S2 is to provide PFC and other frontal areas with sensory information about f2. Here we focus on this initial purely sensory period after the onset of f2. To minimize contamination with decision-related activity, we confine our analyses to responses elicited within the first 181 ms after f2 onset.

Figure 3 shows firing rates in four-example S2 neurons that demonstrate combinations of frequency and context dependence found in the neurophysiological data as a function of the frequency of the vibrotactile stimulus being applied. S2 neurons may or may not exhibit frequency dependence (compare *A*, *C* with *B*, *D*) and context dependence (compare *A*, *B* with *C*, *D*). Moreover, the f2 response of context-dependent neurons may be higher or lower than the f1 response (compare *C* with *D*). Summing over the counts for both monkeys in Table 1, 39% of the S2 neurons are frequency dependent and 12% are context dependent, a small but significant, and as we shall show, sufficient number.

Any context-dependent firing rate shift in our linear model of Equation 1 can be accomplished with a combination of a subtractive (or additive) and divisive (or multiplicative) gain. A subtractive effect corresponds to a decrease in β_2 and a divisive effect to a decrease in β_3 . From Figure 4, we see that context shifts in S2 neurons are overwhelmingly additive or subtractive, although there is also a small and well correlated divisive or multiplicative effect.

As mentioned above, Salinas (2004b) showed that a simple single-layer network model is sufficient to implement drastic changes in functional connectivity, enabling context switching without rewiring. We combine the insights of Salinas' model with

Figure 3. Means and SEs in firing rate during the first 181 ms with respect to the first (f1) and second (f2) stimulus frequencies for four S2 neurons. Neurons may exhibit firing rates that depend on context (C, D), frequency (B, D), or neither (A).

Figure 4. Context-dependent gain modulation of firing rate in plus and minus S2 neurons in terms of the parameters of Equation 1, specifically as the change in slope (β_3 , multiplicative) with respect to the change in intercept (β_2 , additive). The best fit for plus neurons is given by $\beta_3 = -0.0122 - 0.0327 * \beta_2$, and the best fit for minus neurons is $\beta_3 = 0.0025 - 0.0307 * \beta_2$.

the context-dependent neurophysiological data from Romo's laboratory to develop a biologically plausible model of context switching that we then apply to the signal transformation from S2 to PFC.

The simple static network in Figure 5 can solve the sign change problem by taking full advantage of the variety in firing rate modulation observed in S2, thus providing an functional explanation for context-dependent patterns of modulation.

For each stimulus frequency, the inputs to the one-layer network are experimentally observed S2 firing rates, and the output is a weighted sum of the inputs described by Equation 9. Our goal is to find appropriate weights W to transform the data-based inputs x into the desired outputs y , thus implementing the sign switching in stimulus dependence required by earlier decision-making model of our laboratory and observed in PFC (Machens et al., 2005).

We demonstrate the approach with a pair of related examples. For S2 inputs to the network, there are 268 task-responsive neu-

rons that were recorded during 10 pairs of f1 and f2 frequencies, in which both came from the set 10, 14, 18, 22, 26, 30, and 34 Hz. We pretend that all 268 were simultaneously recorded from a single monkey and add a bias term for a total of 269 inputs to the network. We choose outputs to implement an idealized plus input neuron to PFC (y^+) that performs the required sign switch: it reports $44 - f1$ during the first stimulus (f1), and it reports $44 - f2$ during the second stimulus (f2) so that the outputs are 34, 30, ..., 10 for stimulus frequencies of 10, 14, ..., 34. Finally, we calculate the weights based on inputs in two ways: with mean firing rates and with a sampling method. The treatment of minus input neurons to PFC is identical to that of plus neurons, except that the network would report $44 - f1$ during f1 and f2 during f2, so we omit it here. The subsequent analysis will therefore only discuss plus neurons.

In the first case, using mean firing rates for inputs, there are 269 unknown synaptic weights and only 14 constraints, one for each f1 and f2, in the matrix Equation 8. Our mean-based problem is underconstrained, so there is no unique solution for the synaptic weights W . We take a standard approach and use singular value decomposition (SVD) to provide a single solution that minimizes the sum of squared weights.

SVD generates an exact solution for the full mean-firing-rate-based network, with many small (both positive and negative) weights. To measure network robustness, we selectively "prune" connections; that is, we set their weights to zero. Figure 6 compares three cases of pruning with the base case of all weights present. Compared with the perfect matching of desired to actual outputs when all weights are present, eliminating the smallest (in magnitude) half of weights still gives good network performance. Similarly, randomly deleting 10% of weights often results in pretty good performance. However, when half the weights are randomly pruned, performance suffers.

Our first network introduced our model with a simple version based on firing rates averaged over trials. In reality, experiments proceed in a series of individual trials and biological neurons are noisy, so we want to incorporate single-trial information into our synaptic weights. Our model, like the monkeys, must successfully cope with highly variable neuronal firing rates.

This time, instead of averaging firing rates, we take data directly from trials. We want to simulate many single trials from simultaneously recorded single neurons and do so with a sampling procedure. We begin by pretending, as before, that our 268 S2 neurons were simultaneously recorded from a single monkey. For each combination of neuron, context, and frequency, we partition the data into two sets: two-thirds of the experimental trials are randomly assigned to the "training set" and the remaining one-third is assigned to the "test set." As in artificial neural networks, networks determine their weights from the data in the training set, whereas the test set provides an independent measure of performance.

A single simulated trial is a vector whose elements are drawn from the appropriate set of trials of a neuron, plus a bias term. Training trials are sourced from training sets, and test trials are

sourced from test sets. We simulated 4200 training trials, 300 for each context and stimulus, and 420 test trials, 30 for each context and stimulus. SVD on the training data returns the least-squares error-minimizing solution to the weights.

Figure 7 shows that, as expected, the network trained on trial samples is better at dealing with the variability in the test data than the network trained on mean data. In the sample-based network, unreliable neurons tend to be assigned weights close to zero. In contrast, the mean-based network treats all neurons as equally reliable. As a consequence, when presented with data that has trial-to-trial variability, the mean-based network has higher output variance than the trial-based network. Nonetheless, the trial-based responses are still quite variable and often diverge significantly from the desired output. Exactly how much better is the performance of the sample-based network?

We presented both networks with test data on the set of 10 f1–f2 pairs used in the Romo laboratory experiments. For each f1–f2 pair, summing the network output in response to f1, $y^+ = f1$, and the output to f2, $y^- = 44 - f2$, gives $f1 + 44 - f2$, from which we can derive the sign of f1–f2 by a comparison with a threshold. If the estimate of the sign of the difference matches the true sign of f1–f2, we score the answer of our network as correct.

When trained on the full set of neurons, the sample-based network correctly reports the comparison of two stimuli 87% of the time, whereas the monkeys achieve an accuracy of 92%. By comparison, the mean-based network is right only 75% of the time. If instead we base a network on only those 36 neurons that are significantly context dependent, the result is only 65% accuracy. Figure 8 illustrates why. Context dependence is effectively independent of frequency dependence so, to perform well, the neurons in a network must encode both the context and the stimulus frequency accurately over the full range of frequencies. Nonetheless, if we control for neuronal population size, context dependence is important for accuracy. A network consisting of the 134 most context-dependent neurons performs better than one consisting of the 134 least context dependent, at 79% accuracy compared with 72%. Figure 9 summarizes network types and their accuracy. Thus, we need sufficient numbers of both context-dependent and frequency-dependent neurons for good performance. If we further concentrate on the more context-dependent neurons that are also more frequency dependent, as defined by possessing f1 or f2 p values below the median, (70 neurons, approximately one-quarter, for both f1 and f2), performance declines to the level of the less context-dependent network. Neurons with little or

Figure 5. A simple linear feedforward S2 to PFC network solves the context-switching problem. Neither synaptic weights nor neurons change between f1 and f2. x_i indicates the experimentally recorded firing rate of neuron i during stimulus f1; $\tilde{x}_i$ indicates the experimentally recorded firing rate of neuron i during stimulus f2.

Figure 6. Model performance after pruning of selected weights. Percentages in the figure refer to weights remaining after deletion. **A**, No pruning (full set); **B**, the smallest 50% in magnitude are removed; **C**, a random 10% are deleted; **D**, a random 50% are deleted. Deleted weights are set to zero. The top row has histograms of synaptic weights, and the bottom row shows actual network outputs (dots) and desired outputs (lines) for f1 (red) and f2 (blue). **C** and **D** are typical examples.

no context or frequency information still contribute to the performance of the network.

With only 268 neurons, our network comes close to matching the monkeys' performance. Because we are working from a pre-existing dataset, we are unable to increase its size in the near term. Nonetheless, our accuracy shortfall of only 5% suggests that a small increase in the number of input neurons would be sufficient to give us a robust network based on real data that accu-

Figure 7. Outputs of two different networks with naive test inputs. Each point represents a single output corresponding to a single test trial. Network 1 is based on firing rates averaged over trials. Network 2 is based on firing rates from single trials in the form of 4200 samples from each S2 neuron. *A* shows responses of both networks to 210 simulated f1 input vectors, 30 for each frequency. *B* shows responses of both networks to 210 simulated f2 input vectors, 30 for each frequency. Response markers are slightly offset along the x-axis for ease of viewing. Solid lines indicate the desired output. Network 1 has a sample SE of 14.0 compared with a sample SE of 10.2 for network 2.

Figure 8. *p* values for frequency dependence with respect to the first and second stimuli (f1, f2) versus *p* values for context dependence for 268 S2 neurons.

rately performs the computation required for the two-stimulus-interval discrimination task.

Discussion

We have achieved our two aims: (1) to measure the level of context dependence in the secondary somatosensory cortex during a two-stimulus-interval discrimination task, and (2) to use the experimentally observed context dependence as the basis for a simple static network that performs a required context-dependent transformation of stimuli without rewiring.

The low but significant proportion of context-dependent S2 neurons was sufficient to implement the sign inversion required by our recent network model of short-term memory and decision-making (Machens et al., 2005) and, with only 268 neu-

rons plus a bias term, came very close to matching the monkeys' performance. Our results in Figure 9 suggest that a few hundred neurons would be sufficient to match average monkey performance in the two-stimulus-interval vibrotactile discrimination task but remain robust to experimentally observed levels of noise.

Moreover, there are a couple of reasons to think that we may have underestimated the level of context dependence. The first is that neurons were recorded over only a few trials of each f1–f2 stimulus pair, usually resulting in only 5–10 trials of each stimulus. In combination with the strong presence of noise in the data, the 5% significance level in our context hypothesis tests is likely to contribute a few false negatives.

A second possible source of context underestimation is our choice to focus on the number of spikes in the first 181 ms of stimulus presentation. In S2, neuronal activity during f2 initially reflects the purely sensory value of f2 but shifts to depend on (f2 – f1) after 181 ms (Romo et al., 2004). Concentrating on the first 181 ms minimizes contamination by decision-related activity, but a longer period lessens the effect of noise. For example, an increase in the interval to 250 ms, the stimulation

time for required for accurate discrimination during our task (Hernández et al., 1997), increases the number of context-dependent neurons reported by our hypothesis test from 111 to 227 of 922 or from 12.0 to 24.6%.

We thus propose that the monkeys' behavior on this two-stimulus-interval discrimination task (Fig. 1) might be based on inputs provided by $\sim 1/8$ of the neuronal population of area S2. Although this might seem a small fraction, it is not implausible, particularly given the long training times (a few months) required for the monkeys to learn to perform the task. Some of this time might be devoted to selecting context-dependent S2 neurons to preferentially provide inputs to PFC. A testable prediction that follows from this hypothesis is that S2 neurons that show context dependence should be more tightly linked to the animal's behavior than neurons that do not. Thus, choice probability (Britten et al., 1996) should be correlated with context dependence.

For each of our example networks, we could and did choose the outputs to be firing rates with positive or negative unit slopes with respect to stimulus frequency and context. More importantly, our model can transform linear inputs to match any cell with an arbitrary response slope, and, with a sufficient number of nonlinear and context-dependent inputs, it can generate arbitrary nonlinear responses to arbitrary contexts.

An alternative approach to solving the signal transformation problem has been to use an integral feedback signal in PFC to provide an inhibitory signal to upstream neurons (Miller and Wang, 2006). In contrast to our model that takes advantage of the context dependence in S2, the Miller-Wang model (Miller and Wang, 2006) must avoid it, because the model relies on the first stimulus stored in working memory. A context-shifted f2 would generate errors.

In cortex, context dependence is a common phenomenon.

Figure 9. Mean and SE of the accuracy of networks trained on data from neurons with different levels of context dependence and frequency (f1 and f2) dependence. Numbers in brackets indicate the number of neurons in the network before the addition of a bias term. Networks were tested on novel data, and their outputs were judged on the accuracy of their answer to the sign of f1–f2. full set, The network relies on 268 S2 neurons; more context, the 134 neurons more context dependent than the median; less context, the 134 that are less context dependent than the median; context and f1, the “more-context” neurons that are also more f1 dependent than the median (70 neurons); context and f2, defined similarly to “context and f1” (also 70 neurons); significant context, neurons that are context dependent at a significance level of 5% (36 neurons); mean-based, the network that trained only mean firing rates only (from all 268 neurons). The horizontal line represents the mean experimentally observed accuracy.

Neuronal firing rates in many sensory cortices often depend on more than purely sensory features, and the nonsensory modulation has functional consequences. Attention can be thought of as context with respect to saliency and is commonly observed to enhance firing in response to attended stimuli but suppress firing in response to unattended stimuli, in the somatosensory cortex (Hsiao et al., 1993; Burton et al., 1999; Chapman and Meftah, 2005), as well as many other sensory areas. Attentional modulation can improve discrimination (Desimone and Duncan, 1995; Knudsen, 2007), resolve conflicts between incompatible responses (Egner and Hirsch, 2005), increase the probability of making a choice (Lee et al., 1999), and perform coordinate transformations to maintain invariant representations (Salinas and Abbott, 1997). In vision, bistable percepts and binocular rivalry (Leopold and Logothetis, 1996; Maier et al., 2007) are reflected in neuronal responses and can be thought of as a time context. Firing in mitral cells in the olfactory bulb of the rat depend (reversibly) on the context of the predictive value of an odor when it is paired with a reinforcer (whether positive or negative) (Kay and Laurent, 1999).

What are possible mechanisms of context dependence? If we

consider context effects in terms of divisive (or multiplicative, depending on the perspective) and subtractive (or additive) gains, there are many. Divisive effects have been ascribed to a variety, including excitation and inhibition (Murphy and Miller, 2003), active dendrites (Mehaffey et al., 2005), and background synaptic input (Chance et al., 2002). Subtractive effects, in contrast, are usually suggested to arise as a result of shunting inhibition (Holt and Koch, 1997), although the addition of noise and somatic depolarization can also generate divisive effects (Prescott and De Koninck, 2003). In S2, attentional effects have been measured to be primarily multiplicative (Sripati and Johnson, 2006), although we showed in Figure 4 that context-dependent effects in our discrimination task are primarily additive or subtractive. We hypothesize that context-dependent firing is separate from attention and, moreover, given the high correlation between subtractive and multiplicative effects, attributable to a single mechanism.

Salinas (2004b) had proposed a model in which a context- and stimulus-dependent population of sensory neurons could quickly switch between arbitrary desired functions of context and postulated the existence of those sensory neurons. For our part, we needed a mechanism to functionally switch sign as a function of the context of stimulus identity to explain experimentally observed signal transformations between the secondary somatosensory cortex and the prefrontal cortex during a two-stimulus-interval discrimination task, as well as to implement a functional switch in connections as required but not addressed by the recently proposed model by our laboratory of prefrontal cortex processing (Machens et al., 2005). The combination of Salinas’ central idea plus the availability of neurophysiological data recorded by Romo’s laboratory has allowed us to show that context-dependent modulation of functional connectivity is a feasible, biologically plausible and robust solution to a behavioral switching problem in a commonly used two-stimulus-interval discrimination task.

References

- Britten KH, Newsome WT, Shadlen MN, Celebriani S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque mt. *Vis Neurosci* 13:87–100.
- Brody CD, Hernández A, Zainos A, Romo R (2003) Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb Cortex* 13:1196–1207.
- Burton H, Abend NS, MacLeod AM, Sinclair RJ, Snyder AZ, Raichle ME (1999) Tactile attention tasks enhance activation in somatosensory regions of parietal cortex: a positron emission tomography study. *Cereb Cortex* 9:662–674.
- Chance FS, Abbott LF, Reyes AD (2002) Gain modulation from background synaptic input. *Neuron* 35:773–782.
- Chapman CE, Meftah el-M (2005) Independent controls of attentional influences in primary and secondary somatosensory cortex. *J Neurophysiol* 94:4094–4107.
- Courtney SM, Petit L, Haxby JV, Ungerleider LG (1998) The role of prefrontal cortex in working memory: examining the contents of consciousness. *Phil Trans R Soc Lond B Biol Sci* 353:1819–1828.
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222.
- Egner T, Hirsch J (2005) Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nat Neurosci* 8:1784–1790.
- Hernández A, Salinas E, García R, Romo R (1997) Discrimination in the sense of flutter: new psychophysical measurements in monkeys. *J Neurosci* 17:6391–6400.
- Holt GR, Koch C (1997) Shunting inhibition does not have a divisive effect on firing rates. *Neural Comput* 9:1001–1013.
- Hsiao SS, O’Shaughnessy DM, Johnson KO (1993) Effects of selective attention on spatial form processing in monkey primary and secondary somatosensory cortex. *J Neurophysiol* 70:444–447.

- Kay LM, Laurent G (1999) Odor- and context-dependent modulation of mitral cell activity in behaving rats. *Nat Neurosci* 2:1003–1009.
- Knudsen EI (2007) Fundamental components of attention. *Annu Rev Neurosci* 30:57–78.
- Lee DK, Itti L, Koch C, Braun J (1999) Attention activates winner-take-all competition among visual filters. *Nat Neurosci* 2:375–381.
- Leopold DA, Logothetis NK (1996) Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* 379:549–553.
- Machens CK, Romo R, Brody CD (2005) Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* 307:1121–1124.
- Maier A, Logothetis NK, Leopold DA (2007) Context-dependent perceptual modulation of single neurons in primate visual cortex. *Proc Natl Acad Sci U S A* 104:5620–5625.
- Mehaffey WH, Doiron B, Maler L, Turner RW (2005) Deterministic multiplicative gain control with active dendrites. *J Neurosci* 25:9968–9977.
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202.
- Miller P, Wang XJ (2006) Inhibitory control by an integral feedback signal in prefrontal cortex: a model of discrimination between sequential stimuli. *Proc Natl Acad Sci U S A* 103:201–206.
- Murphy BK, Miller KD (2003) Multiplicative gain changes are induced by excitation or inhibition alone. *J Neurosci* 23:10040–10051.
- Prescott SA, De Koninck Y (2003) Gain control of firing rate by shunting inhibition: roles of synaptic noise and dendritic saturation. *Proc Natl Acad Sci U S A* 100:2076–2081.
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP (1992) Numerical recipes in C, Ed 2. Cambridge, UK: Cambridge UP.
- Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4:203–218.
- Romo R, Hernández A, Zainos A, Salinas E (1998) Somatosensory discrimination based on cortical microstimulation. *Nature* 392:387–390.
- Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399:470–473.
- Romo R, Hernández A, Zainos A, Lemus L, Brody CD (2002) Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat Neurosci* 5:1217–1225.
- Romo R, Hernández A, Zainos A (2004) Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron* 41:165–173.
- Salinas E (2004a) Context-dependent selection of visuomotor maps. *BMC Neurosci* 5:47.
- Salinas E (2004b) Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. *J Neurosci* 24:1113–1118.
- Salinas E, Abbott LF (1997) Invariant visual responses from attentional gain fields. *J Neurophysiol* 77:3267–3272.
- Salinas E, Hernandez A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *J Neurosci* 20:5503–5515.
- Sripati AP, Johnson KO (2006) Dynamic gain changes during attentional modulation. *Neural Comput* 18:1847–1867.