A NEURAL ARCHITECTURE FOR RULE-GUIDED BEHAVIOR: A SIMULATION OF PHYSIOLOGICAL EXPERIMENTS

Tetsuto MINAMI¹⁾ and Toshio INUI¹⁾

¹⁾*Kyoto University, Japan*

The ability to use abstract rules is very important in applying our experience to new situations. To examine the system governing such rule-guided behavior, we developed a recurrent neural network model of rule-guided behavior and simulated a physiological experiment involving a rule-guided delayed task (Wallis, Anderson, & Miller, 2001). Our model was constructed using neural system identification (Zipser, 1992), and a fully recurrent neural network model was optimized to perform a rule-guided delayed task. The model's hidden layer contained rule-selective units, and an examination of connection weights substantiated the postulate that rule-selective neurons maintain encoded rule information and contribute to rule-guided responses indirectly. The simulation results predicted functional interactions among neurons involved in various task-related activities. The similarity between the behavior of the model units and biological neurons shows that the brain uses mechanisms like those of the model, and that ample mutual connections in the prefrontal cortex are the basis for promoting flexible learning.

Key words: rule-guided behavior, recurrent network, prefrontal cortex, working memory

The ability to use abstract rules is very important in applying our experience to new situations. Behavior-guiding rules decide which one of different possible outputs will be produced for a given input when more than one response is possible. Rule learning depends on forming associations between disparate, but behaviorally related, information.

Many researchers have argued that rule learning is one of the main functions of the prefrontal (PF) cortex (Passingham, 1993; White & Wise, 1999; Wallis, Anderson, & Miller, 2001; Miller & Cohen, 2001). The WCST (Wisconsin Card Sorting Test) provides evidence connecting the PF cortex with behavior-guiding rules. In the WCST, the subject is asked to match test cards to reference cards according to the color, shape, or number of stimuli on the cards and the sorting rule is changed periodically. The subjects must identify the matching card, learning by trial and error which is the relevant dimension to which they must attend. People with PF cortex damage show stereotypical deficits when performing the WCST (Milner, 1963). That is, they can learn the first rule but then they

This research was supported by the Advanced and Innovational Research Program in Life Sciences, the Project on Neuroinformatics Research in Vision through Special Coordination Funds for Promoting Science and Technology from the Ministry of Education, Culture, Sports, Science, and Technology, of Japan, and a Grant-in-Aid for JSPS Fellows, 14002026, from the Ministry of Education, Science, Sports, and Culture of Japan.

Correspondence concerning this article should be addressed to Tetsuto Minami, Department of Intelligence Science and Technology, Graduate School of Informatics, Kyoto University, Sakyo-ku, Kyoto 606-8501, Japan (e-mail: minami@cog.ist.i.kyoto-u.ac.jp)

cannot escape it: they make a great deal of errors because they perseverate to the earlier rule. Thus, PF cortex damage produces deficits in the application of behavior-guiding rules.

The PF cortex has the ideal infrastructure for a role in behavior-guiding rules. Many PF areas receive converging inputs from at least two sensory modalities, and there are ample interconnections between different PF areas. The PF cortex might select and integrate information from the sensory areas. Several studies have shown that the PF cortex formed the integrative representations of many different types of information, including information about an object and its location (Rao, Rainer, & Miller, 1997; Rainer, Asaad, & Miller, 1998), as well as visual and auditory stimuli (Fuster, Bodner, & Kroger, 2000). Behavior-guiding rules are a more abstract aspect of the integrative representations. The PF cortex activity also seems to represent this information.

Wallis et al. (2001) showed that the neuronal activities of the PF cortex are related to behavior-guiding rules. They explored the neural basis of rule-guided behavior in primates by recording single neurons in the PF cortex of monkeys trained to use two rules: a 'match' rule and a 'non-match' rule. The 'match' rule required each monkey to release a lever if two successive sample objects were identical, whereas the 'non-match' rule required release if the two objects differed. The rule to apply in each trial was indicated by a cue that was presented with the sample object. They found rule-selective neurons that exhibited greater activity when a rule was cued. Other studies have yielded similar results (White & Wise, 1999; Hoshi, Shima, & Tanji, 2000; Asaad, Rainer, & Miller, 2000).

Several reports have shown that the principal function of the PF cortex involves the application of behavior-guiding rules (Wise, Murray, & Gerfen, 1996; White & Wise, 1999). However, the mechanism of this rule-guided behavior is still unknown, and the functional role of rule-selective neurons has not been elucidated. In addition, certain stimulus-selective neurons (rule-selective, object-selective, and multi-selective neurons) must interact with each other in the PF cortex, and the interaction among the neurons affects the rule-guided behavior. Elucidating this functional interaction should lead to understanding of this higher cognitive function.

In this study, we developed a neural network model of a rule-guided delayed match-to-sample task. We were especially interested in how rule-selective neurons contribute to rule-guided behavior. To investigate how the brain may implement a rule-guided behavior, we used recurrent neural networks, using an approach called neural identification (Zipser, 1992). We simulated the physiological results of Wallis et al. (2001), analyzed the temporal patterns of the model units and connection weights, and compared the properties of the units with those of biological neurons.

SIMULATION

The task that we modeled begins with a picture that must be remembered. Simultaneously, a cue is presented that indicates the rule to use. After a delay, the test picture is presented and the monkeys must respond. In match-rule mode, if the pictures

Fig. 1. This is a diagram of the task that we modeled (Wallis et al., 2001). The 'match' rule required each monkey to release a lever if two successive sample objects were identical, whereas the 'non-match' rule required release if the two objects differed. In each trial, a cue presented with the sample object indicated the rule to use.

are the same, the monkeys release the lever; if the pictures differ, they continue to hold the lever. In non-match-rule mode, the reverse is true; i.e., if the pictures differ, the monkeys release the lever, and if they are the same, they hold the lever. A typical trial of the task is presented in Fig. 1.

To complete this task, the network must map from the cue inputs to rule information, store this rule information and cued object information, compare the current object information with stored information, and produce output according to the current rule mode.

The model suggested here was generated using a neural system identification technique (Zipser, 1992). This technique requires specifying the temporal sequence of the input-output behavior rather than the detailed internal structure of the model. The internal mechanism by which input is converted to output is generated indirectly by the application of an optimization paradigm such as gradient descent. Neural system identification is interesting as a modeling paradigm because of the empirical finding that the internal behavior of neural identification models often mimics the internal behavior of the neural plant being modeled.

The models use simple model neurons with a logistic function. Logistic units represent the average spiking rate of real neurons with a continuous variable, and the synaptic input as a weighted sum of the activations of other neurons. These weights are generally chosen automatically using some optimization procedure, which is simply a mathematical technique that chooses the best set of weights, and is unlikely to mimic biological processes.

A recurrent network with feedback connections is interesting for various reasons. A biological neural network has highly recurrent connections. One characteristic of a fully

Fig. 2. The concept underlying the model. We assumed that the PF area is one uniform layer. We regarded the inferior temporal cortex and posterior parietal cortex as the input, and the motor-related area as the output.

recurrent network is that it can retain information for some time using a fixed point attractor. Zipser (1991) showed that the results of a physiological experiment examining neuronal activity during a working memory task were explained fully using the dynamics of this recurrent network. Much research has used this neural identification paradigm (Zipser & Andersen, 1988; Zipser, 1991; Zipser, Kehoe, Littlewort, & Fuster, 1993; Moody, Wise, di Pellegrino, & Zipser, 1998; Moody & Wise, 2000; Xing & Andersen, 2000; Mitchell & Zipser, 2001).

An important aspect of the model is that the neuronal representation in the PF cortex is optimized through learning. As shown previously, the PF cortex is connected with virtually all sensory neocortical and motor systems, as illustrated in Fig. 2. According to the physiological results (Wallis et al., 2001), in the PF area there is little difference in the distribution of neurons preferring different types of information. Moreover, another study (White & Wise, 1999) found no significant regional segregation of such selectivity. Therefore, we assumed that the PF area is one uniform layer. This assumption is not implausible from a biological perspective. The model is a fully recurrent neural network model whose neurons are mutually connected.

The model architecture is shown in Fig. 3. It has input, hidden, and output layers. We have already proposed a neural network model of working memory in order to shed light on its retention mechanism and simulated the physiological results of Rao et al. (1997) (Minami & Inui, 2001). However, previous models contained comparator units outside the hidden layer. The current model integrated this function into the hidden layer.

The input layer consists of object inputs and cue inputs. The input representation that we adopted was chosen to give the information required for the task. The object input was implemented with four discrete binary parameters (e.g., (0, 1, 0, 0)) representing the sample object stimulus. The cue inputs consisted of four inputs representing cue stimuli (C_1 , C_2 , C_3 and C_4). C_1 and C_2 indicate application of a match-rule and C_3 and C_4 indicate

Fig. 3. The neural network architecture includes input, hidden, and output layers. The input layer consists of an object input and a cue input. The hidden layer consists of recurrently connected logistic units.

a non-match-rule. Thus, we used four cue inputs for two rules, because two distinct cues from different modalities were used to indicate the same rule in Wallis's experiment. The network output consists of two units representing the 'response': output (1, 0) meant 'release the lever', while (0, 1) meant 'continue holding the lever (no response)' in the physiological experiments. The hidden layer consisted of recurrently connected logistic units. The network model contained between 20 and 40 of the hidden units. The training of the minimal network was apt to fail. The model of more hidden units was apt to converge faster, but have more dead units that didn't participate in the processing. So, we adopted the model of 30 hidden units as the main target of the analysis.

The output of the i th model unit in the hidden layer on time step $t+1$ is given by

$$h_i(t+1) = f(\sum_j h w_{ij} h_j(t) + \sum_k v_{ik} z_k(t) - b_i) \quad (1)$$

Fig. 4. Schematic of the training algorithm used for the network. The left panel shows the series of cue inputs and the middle panel shows those of sample object inputs. The right panel shows the desired outputs.

where hw_{ij} represents recurrent connections from every hidden unit and v_{ik} represents weighted connections from every input unit z_k . Each unit has a bias, b_i . $f(x)$ is the sigmoid function

$$f(x) = \frac{1}{1 + e^{-x}} \quad (2)$$

Each motor output on time step $t+2$ is defined as

$$o_m(t+2) = f\left(\sum_i ow_{mi} h_i(t+1) - b_m\right) \quad (3)$$

where ow_{mi} represents the weighted connections from the hidden units. Each unit has a bias, b_m .

The connections are optimized through a gradient descent using real-time recurrent learning (RTRL) (Williams & Zipser, 1989). Training started from randomized connection weights. For training, models were generated for roughly 10^7 network time steps. Training ceased when the mean square error was less than .03.

To train the model, randomized sequences of object and cue inputs were provided at appropriate times. In the match-rule mode, the output R_1 was trained to produce a response when the current object input matched the previously presented object input, whereas in the non-match-rule mode, the output R_2 was trained to produce a response when the current object input did not match the previously presented object input. Fig. 4 provides details of the training paradigm.

RESULTS

To determine the mechanism of the model, several analytical approaches were used, based on the following points: the activity pattern during each trial period, the

Fig. 5. Tuning curves for four units that showed selectivity to rules and objects during the a) cue and b) delay periods. The activity is plotted as a function of the cue (C_1, C_2, C_3, C_4) and sample object input (S_1, S_2, S_3, S_4).

consequences of single-unit lesions, the synaptic weight connections, and comparisons with physiological data.

Fig. 6. Typical temporal behavior of a rule-selective unit. The graph shows the temporal pattern when each sample object (S_1, S_2, S_3, S_4) and each cue (C_1, C_2, C_3, C_4) was input at step 5.

Activity patterns during each trial period

We run the network in 10 discrete time steps for each input pattern. In a typical simulated behavioral trial, cue and sample object information were input at step 5, which corresponded to the cue period. After a delay (from steps 6 to 9), an object was input at step 10, which corresponded to the response period. We investigated the temporal patterns of the hidden units in a typical model network consisting of 30 hidden units. The spectrum of the pattern of hidden unit activity ranged from sustained activity during the delay to elevated activity only at the time of the input. We examined the selectivity of the units during the cue, delay, and response periods.

During the cue and delay periods, many hidden units showed rule and object selectivity. We called these units rule-selective and object-selective units, respectively. Fig. 5 shows a cue-tuning curve and an object-tuning curve during the cue and delay periods of a rule-selective unit and an object-selective unit. The rule-selective units showed greater responses to C_1 and C_2 , indicating the match-rule regardless of which cue signified the rule or which object was remembered. Conversely, the object-selective units showed selectivity to stored object information, regardless of which cue was input; they

Fig. 7. Response unit activity during the response period when storing different sample objects. Each plot shows the same model unit when it stored four different sample objects (S_1 , S_2 , S_3 , S_4). The activity is plotted as a function of the current object input when each cue (C_1 , C_2 , C_3 , C_4) was input.

were also similar to the object-tuned units of Minami and Inui (2001). In a typical network consisting of 30 hidden units, 7–10% of the units had selectivity for the learned rules and 20–30% had selectivity for the stored object information. In addition, mixed type units showed selectivity to both cue and object information. These mixed types had activity profiles that seemed to combine rule and object selectivity, to various degrees. Such units were more likely to be observed in smaller networks (e.g., a network consisting of 20 units).

Fig. 6 shows typical temporal patterns of rule-selective units. The rule-selective units showed greater activity during the delay period in match trials, regardless of which cue signified the rule or which object was remembered. The units exhibited the rule selectivity shown in Wallis et al. (2001), which contributed to retaining the two learned rules.

In addition, during the response period, many units showed selectivity to rule and object information. Moreover, some units showed selectivity to both current input and the stored object. We called these units “response units”. In the response units, the stored sample object modulates the amplitude of the response to current stimuli. This

modulation includes both potentiation and suppression. An example of how these units respond to all the current stimuli for various stored sample object information is shown in Fig. 7. These response units have dual-mode activities (match-rule mode and non match-rule mode), as shown in Fig. 7. The shape and peak of the tuning curve remain fixed for different stored objects. The magnitude of the curve is dependent on the current stored objects. Wallis et al. (2001) did not report this type of neuron. Our simulation results predicted that such neurons would be found during the response period in the PF cortex.

Lesion effects of a single unit

In order to explore the role of such rule-selective units in the network, we examined the effects of removing each type of rule-selective unit. Fig. 8 shows the output unit activity for intact (solid line) and lesioned (dashed line) models in both match-rule and non-match-rule mode: a) shows the effect of removing a match-rule-selective unit and b) shows the effect of removing a non-match-rule-selective unit. The dashed lines show the effects of a lesion. A lesion of a match-rule unit affected the output in match-rule mode, but not in non-match-rule mode. The effect includes both potentiation and suppression. Conversely, a lesion of a non-match-rule unit affected the output in match-rule mode. Combined, these data indicate that the two modes (match-rule and non-match-rule mode) operate separately. In addition, Fig. 9 shows the effects of removing an object-selective unit. A lesion of an object-selective unit affected the output, regardless of whether a match-rule or non-match-rule was applied. This suggests that separate rule processing and object processing streams exist in the network. These systems were confirmed in the analysis of connection weights below.

Synaptic weight connections

Next, we investigated the weights of the synaptic connections in the model. An examination of the weights between the input layer and the hidden layer showed that rule-selective units had greater connections with cue inputs, whereas they had little connections with object inputs. Moreover, the units that preferred the match-rule had positive connections with C_1 and C_2 and negative connections with C_3 and C_4 . The units that preferred the non-match-rule behaved in the opposite manner. This shows that rule-selective units contributed to encoding rule information. Object-selective units had relatively large connective weights with object inputs.

As shown by the selectivity during the response period, the examination of the weights between the hidden and output layers showed large weights for synapses connecting response units to the output. We divided these response units into two groups: R_1 and R_2 response units. R_1 response units had large positive weights with R_1 output units and large negative weights with R_2 output. That is, R_1 response units facilitated the R_1 response and suppressed the R_2 response. The converse was true for R_2 response units. Moreover, most rule-selective and object-selective units had weaker direct influences on the output units.

Object-selective units received object input, had a strong recursive connection, and were mutually connected. Similarly, each rule-selective unit had a strong positive

Fig. 8. Effect of removing each type of rule-selective unit: a) the effect of removing a match-rule-selective unit and b) the effect of removing a non-match-rule-selective unit. The top two columns show the output unit activity when a match-rule was applied (when C_1 and C_2 was input), and the bottom two columns show the output unit activity when a non-match-rule was applied (when C_3 and C_4 was input). Each column shows R_1 and R_2 output unit activity for intact (solid line) and lesioned (dashed line) models. Each plot shows the same model unit when it stored four different sample objects (S_1, S_2, S_3, S_4). The activity is plotted as a function of the current object input when each cue was input. Each arrow shows the activity shift resulting from the removal of a unit.

Fig. 9. Effects of removing an object-selective unit. The top two columns show the output unit activity when a match-rule was applied, and the bottom two columns show the output unit activity when a non-match-rule was applied. Each column shows R_1 and R_2 output unit activity for intact (solid line) and lesioned (dashed line) models. Each plot shows the same model unit when it stored four different sample objects (S_1, S_2, S_3, S_4). The activity is plotted as a function of the current object input when each cue was input. Each arrow shows the activity shift with the removal of a unit.

connection with itself. The match-rule-selective units had negative weights with non-match-rule-selective units and vice versa. These findings show that rule-selective units maintain rule information using recurrent connections among rule-selective units. The connection weight from rule-selective units to response units was relatively large, but the opposite connections were very weak, which shows that response units receive rule information from rule-selective units and use this information to judge which mode (match-rule mode or non-match-rule mode) should be applied. Furthermore, response units receive input from both the current stimulus and the retained object units. The two types of response units do not interact very much.

These analyses of connection weights in the network model can help to elucidate the mechanism for performing rule-guided behavior. The network computed output responses in the following way. (1) A rule-selective unit received input from a cue input, encoding and maintaining rule information. (2) An object-selective unit received input from an object input and maintained object information. (3) A response unit received a stored rule, and object information from rule-selective units and object-selective units, compared them with the current inputs, and produced output signals (Fig. 10).

These processing systems are consistent with the results of the lesion study above.

Fig. 10. Schematic of the network mechanism: the relationships among the input layer, hidden units, and the output layer.

These results suggest that the computational solution in our model can be implemented in a network that includes PF cortex neurons.

Comparisons with physiological data

The hidden layer units in our model were assumed to correspond to the neural networks centered on the PF cortex in the brain. We compared the properties of the model units with those that Wallis et al. (2001) obtained for PF cortex neurons in physiological experiments. Our results confirm that rule-selective neurons in the PF cortex encoded and maintained rule information throughout the simulation.

In their physiological experiment, Wallis et al. (2001) did not examine neuronal activity during the response period. Our results predict that some response-related neurons will be found in the brain system centered on the PF cortex if the neuronal activities during the response period are investigated (Fig. 7).

Funahashi and Inoue (2000) examined functional interactions among PF neurons during a simple oculomotor delayed-response task, using a cross-correlation analysis. They diagrammed the potential neuronal circuitry in the PF cortex necessary to perform the task (Funahashi & Inoue, 2000, Fig. 13). Our results predict that diagrams similar to

Fig. 10 would be obtained through a cross-correlation analysis of neuronal activity during rule-guided behavior.

DISCUSSION

Many studies have investigated rule-guided behavior, and neuropsychological and neurophysiological studies have shown that the PF cortex plays a major functional role in such behavior. However, the mechanisms for rule-guided behavior remain unclear. In this study, we investigated rule-guided behavior, proposed a neural network model to shed light on its mechanism, and simulated physiological experiments that examined the coding of rules.

The main feature of our model is that it is a recurrent neural network, which is interesting for various reasons. A biological neural network has highly recurrent connections. One characteristic of a fully recurrent network is that it can retain information for some time using a fixed point attractor. Zipser (1991) showed that the physiological results of neuronal activity during a working memory task can be fully explained using the dynamics of a recurrent network.

The model described here augments previous models in several ways. Although active storage has been modeled (Zipser, 1991), and a matching function incorporated (Moody et al., 1998), these have not been previously combined with a rule-guided function. In addition, our model does not assume a gate or load input with an ambiguous source. The representation of output units is directly linked with motor control, and is plausible.

The rule-selective units had the same properties as the rule-selective neurons shown in Wallis et al. (2001). Analysis of the connection weights shed light on the mechanism by which the PF cortex performs a rule-guided delayed matching-to-sample task, and elucidated the functional role of rule-selective neurons. The finding that the rule-selective units had stronger connections with cue inputs, and were mutually connected, indicates that the units contribute to the coding and retention of rule information. Their removal affects each rule mode, which shows that each type of rule-selective unit affects the matching output separately. Mixed type units seemed to participate in several of the information-processing components to various degrees. The response units showed the response selectivity during the response period, but most of them work as the mixed type units during the delay period, so these units seem to be important for integrating information. In addition, not only the recurrent connections among the same kind of units but also the interaction between the mixed type units and pure functional type units (the rule-selective and object-selective units) contribute to the retention of the information, especially in smaller networks.

One prediction of the model is that neuronal interactions occur in the PF cortex, in a manner similar to the diagrams in Fig. 10. In our simulation, we hypothesized that the hidden layer was a neural network centered on the PF cortex. This hypothesis is consistent with recent neuroanatomical evidence, which indicates a high degree of

intrinsic excitatory connectivity in the PF cortex (Pucak, Levitt, Lund, & Lewis, 1996). It has also been suggested that neurons with similar response preferences in the PF cortex are interconnected (Pucak et al., 1996). This suggestion is not inconsistent with our finding that selective units of the same kind had strong recursive connections.

We simulated the results of a physiology experiment using a recurrent network model with a simple input-and-output relation for executing a rule-guided delayed task. Our model explained the function of the PF cortex under the assumption that the input module corresponds to the inferior temporal cortex and other brain areas, and the output layer to the motor-related area. Of course, it is difficult to directly connect RTRL with the learning mechanisms of the cerebral cortex. However, the close similarity between the behavior of the model units and biological neurons suggests that the brain uses mechanisms like those of the model, and our results suggest that ample mutual connections in the PF cortex are the basis for promoting flexible behavior, such as rule learning. Moreover, the fact that our model simulated several physiological experiments well demonstrates its validity.

In this study, a processing network such as that shown in Fig. 10 emerged automatically by optimizing a fully recurrent network without assuming specific connections. This indicates the flexibility of a recurrent network, and this flexibility is an important feature of the PF cortex. The architecture of our model is very simple, but our model has an advantage over other models for rule-learning (O'Reilly, Noelle, Braver, & Cohen, 2002) in that the PF representations are not hand-coded.

In summary, our simulation demonstrated that rule-selective neurons play a functional role in the coding and retention of rule information, which is used to produce the correct response appropriate to the context. Future work that considers the anticipatory dynamics and more reasonable representations of the cues should contribute to furthering our understanding of rule-guided behavior.

REFERENCES

- Asaad, W. F., Rainer, G., & Miller, E. K. 2000. Task-specific neural activity in the primate prefrontal cortex. *Journal of Neurophysiology*, **84**, 451–459.
- Funahashi, S., & Inoue, M. 2000. Neuronal interactions related to working memory processes in the primate prefrontal cortex revealed by cross-correlation analysis. *Cerebral Cortex*, **6**, 535–551.
- Fuster, J. M., Bodner, M., & Kroger, J. K. 2000. Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature*, **405**, 347–351.
- Hoshi, E., Shima, K., & Tanji, J. 2000. Neuronal activity in the primate prefrontal cortex in the process of motor selection based on two behavioral rules. *Journal of Neurophysiology*, **83**, 2355–2373.
- Miller, E. K., & Cohen, J. D. 2001. An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, **24**, 167–202.
- Milner, B. 1963. Effects of different brain lesions on card sorting. *Archives of Neurology*, **9**, 90–100.
- Minami, T., & Inui, T. 2001. A neural network model of working memory: Processing of “what” and “where” information. In J. Mira & A. Prieto (Eds.), *IWANN. Vol. 2084 of lecture notes in computer science* (pp. 126–133). New York: Springer.
- Mitchell, J., & Zipser, D. 2001. A model of visual-spatial memory across saccades. *Vision Research*, **41**, 1575–1592.
- Moody, S. L., & Wise, S. P. 2000. A model that accounts for activity prior to sensory inputs and responses

- during matching-to-sample tasks. *Journal of Cognitive Neuroscience*, **12**, 429–448.
- Moody, S. L., Wise, S. P., di Pellegrino, G., & Zipser, D. 1998. A model that accounts for activity in primate frontal cortex during a delayed matching-to-sample task. *Journal of Neuroscience*, **18**, 399–410.
- O'Reilly, R. C., Noelle, D. C., Braver, T. S., & Cohen, J. D. 2002. Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. *Cerebral Cortex*, **12**, 246–257.
- Passingham, R. 1993. The frontal lobes and voluntary action. Oxford: Oxford University Press.
- Pucak, M. L., Levitt, J. B., Lund, J. S., & Lewis, D. A. 1996. Patterns of intrinsic and associational circuitry in monkey prefrontal cortex. *Journal of Comparative Neurology*, **376**, 614–630.
- Rainer, G., Asaad, W. F., & Miller, E. K. 1998. Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature*, **393**, 577–579.
- Rao, S. C., Rainer, G., & Miller, E. K. 1997. Integration of what and where in the primate prefrontal cortex. *Science*, **276**, 821–824.
- Wallis, J. D., Anderson, K. C., & Miller, E. K. 2001. Single neurons in the prefrontal cortex encode abstract rules. *Nature*, **411**, 953–956.
- White, I. M., & Wise, S. P. 1999. Rule-dependent neuronal activity in the prefrontal cortex. *Experimental Brain Research*, **126**, 315–335.
- Williams, R. J., & Zipser, D. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, **1**, 270–280.
- Wise, S. P., Murray, E. A., & Gerfen, C. R. 1996. The frontal cortex-basal ganglia system in primates. *Critical Review of Neurobiology*, **10**, 317–356.
- Xing, J., & Andersen, R. A. 2000. Memory activity of lip neurons for sequential eye movements simulated with neural networks. *Journal of Neurophysiology*, **12**, 429–448.
- Zipser, D. 1991. Recurrent network model of the neural mechanism of shortterm active memory. *Neural Computation*, **3**, 179–193.
- Zipser, D. 1992. Identification models in the nervous system. *Neuroscience*, **47**, 853–862.
- Zipser, D., & Andersen, R. A. 1988. A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, **331**, 679–684.
- Zipser, D., Kehoe, B., Littlewort, G., & Fuster, J. 1993. A spiking network model of short-term active memory. *Journal of Neuroscience*, **13**, 3406–3420.

(Manuscript received December 10, 2003; Revision accepted January 13, 2004)