SPATIAL PERCEPTION AND MOVEMENT PLANNING IN THE POSTERIOR PARIETAL CORTEX

by

PIETRO MAZZONI

B.S., Physics, University of California, San Diego (1988)

Submitted to the Department of Brain and Cognitive Sciences in Partial Fulfillment of the Requirements for the Degree of

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 1994

Signature of Author _____
Department of Brain and Cognitive Sciences
January 31st, 1994

Certified by _____
Richard A. Andersen
Professor of Brain and Cognitive Sciences
Thesis Supervisor

Accepted by _____
Emilio Bizzi
Eugene McDermott Professor in the Brain and Cognitive Sciences and Human Behavior
Head, Department of Brain and Cognitive Sciences

SPATIAL PERCEPTION AND MOVEMENT PLANNING IN THE POSTERIOR PARIETAL CORTEX

by

PIETRO MAZZONI

Submitted to the Department of Brain and Cognitive Sciences on January 31st, 1994, in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Neuroscience

ABSTRACT

In this dissertation I describe a series of studies designed to elucidate the role of the posterior parietal cortex (PPC) of the primate brain in spatial perception and movement planning. In these studies my collaborators and I examined these processes in the context of sensorimotor integration, that is, the formulation of movement plans based on specific sensory cues. We chose saccadic eye movements as our model sensorimotor behavior. Saccades are high-speed eye movements that redirect the line of sight from one point of interest in the visual field to another. We used theoretical approaches involving neural network models and experimental techniques consisting of recording the activity of single neurons in the brains of awake behaving monkeys.

Sensorimotor behaviors such as saccadic eye movements require a set of transformations of the incoming signals. One such computation is the transformation of the coordinate frame in which the location of a sensory stimulus is encoded. I first review a series of neural network models that were developed in our laboratory to model how such transformations might be carried out in the PPC (Chapter 2). These models showed that a set of neurons with the properties observed in the PPC can support a distributed code of a stimulus' spatial location in a new coordinate frame, different from the reference frame of the incoming signal. I then relate a study addressing the plausibility of these neural networks as models of a biological structure such as the PPC (Chapter 3). This study showed that more biologically plausible versions of the original neural network can still implement coordinate transformations in the manner suggested by the properties of PPC

neurons, thus strengthening the validity of this modelling approach in the study of PPC function.

Another transformation required for spatial behavior is the encoding of spatial cues obtained though different sensory modalities into a single spatial signal. I first review the role of various nervous system structures in encoding the locations of sound sources (Chapter 4). Next I describe experiments showing that the lateral intraparietal area (area LIP) of the PPC contains neurons encoding the locations of auditory as well as visual stimuli. The responses to auditory stimuli are modulated by eye position in a manner that allows area LIP to encode the locations of auditory and visual signals in a common reference frame (Chapter 6).

The third issue I address is the formation of a motor plan. In one set of experiments we asked whether area LIP is solely concerned with sensory stimuli or whether it expresses any aspect of the movement being planned (Chapter 7). We show that most neurons in this area unambiguously encode the next planned saccade, establishing a direct role of this area in movement planning. This result predicted that if a monkey changed his saccade plan, the activity of its area LIP neurons should reflect such a change of plan even when the monkey makes no movement. When we tested this prediction experimentally we found that the motor planning activity of most LIP neurons can be turned on and off as the animal changes its movement plan, establishing that this activity reflects the monkey's covert intention even in the absence of behavior (Chapter 8).

Finally, I show that temporary inactivation of the PPC produces marked reversible deficits in a monkey's ability to make saccades to the remembered locations of visual stimuli, while producing only small deficits in saccades performed under direct visual guidance (Chapter 9). These results suggest an essential role of the PPC in the programming of saccades that are more difficult or that require additional processing compared to ordinary visual saccades.

Thesis Supervisor: Richard A. Andersen
Title: Professor of Brain and Cognitive Sciences

ACKNOWLEDGEMENTS

I have had a wonderful time while a graduate student at M.I.T. I would not be able to say so had the following people not been around me to help, teach, advise, stimulate, guide, enrich, entertain, humor, and support me throughout these years.

I thank Richard Andersen for his unwavering support, encouragement, and guidance.

I thank Emilio Bizzi, Michael Jordan, and James Lackner for kindly serving on my thesis committee.

I thank Catherine Cooper, Gail Robertson, Steve Marchetti, and Richard Keough for constant, invaluable administrative and technical assistance.

I thank Martyn Bracewell, Shabtai Barash, Michael Jordan, and Brigitte Stricanne, who put up with me as a collaborator and still became my friends.

I thank Ray Li for the unlimited cooperativeness and kindness he demonstrated while sharing the experimental set-up with me.

I thank Stefan Treue for teaching me more about Macintosh computers than I had thought I ever wanted to know; David Bradley for introducing me to more statistical techniques than may be safe in my hands; Peter Thier for teaching me several experimental techniques; Ning Qian and Larry Snyder for their invaluable help with UNIX; and Bard Geesaman for widening my musical repertoire (which preserved a shred of my sanity during months of data analysis).

I thank all current and previous members of the Andersen lab (including, besides those mentioned above, Peter Brotchie, Sabrina Goodman, Michael Graziano, Bob Snowden, Leo Fogassi, and Jing Xing) for always being forthcoming with technical advice, practical assistance, stimulating discussions, and moral support. The presence of these people created an atmosphere in the lab that made it very enjoyable to carry out research in it.

I especially thank Jan Ellertsen for guiding my progress and keeping me on a steady forward course in the Department of Brain and Cognitive Sciences at MIT. I also thank Debbie Reynolds for her practical assistance in departmental matters.

I thank Linda Burnley and her staff at the M.D.-Ph.D. Program of Harvard Medical School, and Keiko Oh, Walter Koltun, Fred Bowman,

Barbara Jaskela, and Patty Cunningham of the Harvard-MIT Division of Health Sciences and Technology for guiding me through—and often rescuing me out of—the mazes of the MIT and Harvard bureaucracies.

I thank Sungsil Kim, Louisa Worthington, and Pat Claffey for their help in using the M.I.T. libraries.

I am grateful for the Dennison Fellowship and the Irene and Charles Albisetti Fellowship, which partly supported me during my studies.

Finally I thank my parents, my sisters, my brother, and Mary Ann for always being near me.

CONTENTS

III. Forming A Motor Plan

BIOGRAPHICAL NOTE

I was born in Rome, Italy, on December 10th, 1966. I grew up in Valmontone, near Rome, until 1982, when I moved with my family to San Diego, California. I studied classical piano from the age of 9, attended the Conservatorio di Musica "Santa Cecilia" in Rome from 1981 to 1982, and majored in music (piano) at the University of California, Los Angeles, from 1983 to 1984. Then I switched to a physics curriculum at the University of California, San Diego, where I obtained a B.S. in 1988. Since then I have been in the M.D.-Ph.D. program at Harvard Medical School in Boston. In 1990 I paused my medical studies and entered the doctoral program of the Department of Brain and Cognitive Sciences at MIT. I have been in the laboratory of Richard Andersen since 1988.

PUBLICATIONS

1. Mazzoni, G., and Mazzoni, P. (1986) Reversibility of graft rejection following a short cyclosporin-A administration 5 days after grafting. *Microsurgery* 7(3):114-116.
2. Mazzoni, G., and Mazzoni, P. (1988) Hyperacute rejection after portacaval shunt in the rat. *Transplantation Proceedings* 20(Suppl.):844-847.
3. Selverston, A. I., and Mazzoni, P. (1989) Flexibility of computational units in invertebrate CPGs. In *The Computing Neuron*, R. M. Durbin, R. C. Miall, and G. J. Mitchison, eds., pp. 205-228, Addison-Wesley, New York.
4. Mazzoni, P., Andersen, R. A., and Jordan, M. I. (1990) A_{R-P} learning applied to a network model of cortical area 7a. *Proceedings of the International Joint Conference On Neural Networks*, San Diego, CA, June 1990, vol. II, 373-379.
5. Mazzoni, P., Andersen, R. A., and Jordan, M. I. (1991) A more biologically plausible learning rule than backpropagation applied to a network model of cortical area 7a. *Cerebral Cortex* 1(4):293-307.

6. Mazzoni, P., Andersen, R. A., and Jordan, M. I. (1991) A more biologically plausible learning rule for neural networks. *Proceedings of the National Academy of Sciences* 88:4433-4437.

7. Andersen, R. A., Brotchie, P. R., and Mazzoni, P. (1992) Evidence for the lateral intraparietal area as the parietal eye field. *Current Opinion in Neurobiology* 2:840-846.

8. Mazzoni, P., and Andersen, R. A. (1993) The role of neural networks in the study of the posterior parietal cortex. *Proceedings of the International Conference On Neural Networks*, San Francisco, CA, March 1993, vol. III, 1321-1326.

9. Mazzoni, P., and Andersen, R. A. (in press) Coordinate transformations for gaze coding in the posterior parietal cortex. In *The Handbook of Brain Theory and Neural Networks*, M. A. Arbib, ed., MIT Press, Cambridge, MA.

A papá e mamma

Chapter 1

Introduction

Humans and many other animals most commonly shift their gaze by making saccadic eye movements. Saccades are high-speed movements that redirect the line of sight from one point of interest in the visual field to another. In order to shift our gaze we must first perceive a location in the visual field, formulate a motor plan, and execute a saccade towards the selected location. Saccades thus require sensorimotor integration, that is, the integration of sensory signals encoding spatial locations and the formulation of movement plans based on these signals.

The posterior parietal lobe of the primate brain has long been considered important for sensorimotor behaviors such as saccades. Specific areas within the posterior parietal cortex (PPC) appear to compute sensorimotor transformations for the purpose of programming the direction of gaze. In this thesis I describe several approaches we used in our laboratory to understand the role of the PPC in the production of saccadic eye movements. I report the results of theoretical and experimental studies addressing the relationship between neural activity in the PPC and some of the transformations of neural signals required to perceive the location of a sensory stimulus and to direct gaze towards it.

One class of operations inherent to any sensorimotor behavior are coordinate transformations. The coordinates of a neural signal can in general refer to any of a number of features encoded by that signal. If a signal encodes a spatial location, then its coordinates are defined by the reference frame relative to which the location is defined. Because the sensory organs of many animals can move relative to other parts of the animal's body, spatial neural signals can be encoded in a variety of coordinate frames. In the primate spatial information obtained via the visual sense, for example, is originally encoded in a reference frame based on the eyes, whereas the physical cues that inform the nervous system of the location of a sound are anchored to a head-centered reference frame. These two reference frames are distinct because the eyes can rotate relative to the head. In order to combine visual and auditory spatial information, therefore, the primate's nervous system

must transform these spatial signals from one coordinate frame to another. If a visual and auditory stimulus such as a hissing snake, for example, appears in front of primate and the primate wants to look at it (before running away), the visual and auditory spatial information must be combined in order to program the correct saccade.

Sensorimotor behaviors such as saccades require other operations besides the transformation of coordinates. One of these is the transformation of sensory signals into motor commands. These presumably involve a number of steps, including but not limited to the transformation of the signals' spatial coordinates. These processes can be complex and loosely defined, and can include the perception of a sensory stimulus' various features, the allocation of attention, interactions of the stimulus' perception with stored memories, and so on, right up to the formation of a set of motor commands that will produce the behavior elicited by the stimulus' presence. Among the processes one that seems—at least intuitively—required for the generation of a motor command based on a sensory signal is motor intention. This process is the expression of a plan to execute a particular movement that precedes the actual command that produces the movement. Every deliberate movement—as opposed to a "reflexive" or "automatic" one—is by definition preceded by the intention to make that movement. One might expect such an intention to be expressed as neural activity somewhere in the brain in relation to at least certain classes of movements.

Neurological and physiological evidence has for long made the PPC of the primate brain a candidate area for participating in the classes of transformations I have described—the transformations of spatial coordinates and the linking of sensory and motor signals. The underlying theme of this thesis is an analysis of whether and how the activity of PPC neurons can subserve these transformations. An overview of how these topics are covered in the thesis' chapters follows.

SECTIONS I-II

The first two sections (**Chapters 2-6**) are devoted to the issue of coordinate transformations in the PPC.

I first review a neural network modelling approach that has been used to study how coordinate transformations can be achieved by the PPC (**Chapter 2**). These models show that visual information and eye position signals interact in at least two areas of the PPC to produce a head-centered code of a stimulus' location distributed over a neuronal population.

A problem with the original neural networks used to model PPC function was their plausibility as models of brain function. These models were trained to compute coordinate transformations using a learning algorithm (backpropagation) considered implausible as a mechanism for biological learning. In **Chapter 3** I relate a theoretical study in which we addressed the issue of the biological plausibility of neural networks as models of brain function. We show that another class of neural network models can learn the same coordinate transformation task and develop the same distributed head-centered code as observed in the PPC and in the other model networks.

Next I address the issue of localization of sound sources. Sensorimotor integration involves not only the transformation of the spatial coordinates of individual visual stimuli. The cues to a spatial location in the environment can take many forms and enter the nervous systems through the visual, auditory, and somatic sensory modalities. We thus asked whether the PPC can process spatial cues presented through the auditory system.

Chapter 4 provides a background of lesion studies that have identified various brain structures as necessary for various aspects of auditory localization. In **Chapter 5** I describe responses of neurons in the lateral intraparietal area (area LIP, a subdivision of the PPC) to auditory stimuli. These responses are spatially tuned, and many of the neurons exhibiting them also have spatially tuned visual responses. Area LIP may thus encode the locations of auditory as well as visual

stimuli. This result begs the question of the coordinate systems in which the auditory signals are encoded and how these coordinates relate to those of the visual signals. Initial results of ongoing experiments examining these questions are reported in **Chapter 6**.

SECTION III

This section (**Chapters 7-9**) is devoted to the distinction between signals related to sensory stimuli and those related to motor preparation, and to the functional role of PPC activity in saccade generation.

Chapters 7 and 8 examine whether area LIP of the PPC contains signals related to the motor preparation of saccades. We show that the activity of most LIP neurons reflects the intention to execute the next planned saccade. The PPC thus not only transforms the spatial coordinates of sensory stimuli, but also processes signals expressing the preparation of movement.

Finally, in **Chapter 9** we ask whether the activity of PPC neurons is necessary for any particular aspects of saccade programming. Temporary inactivation of a portion of the PPC produced changes in saccade metrics that were particularly pronounced for saccades to the locations of remembered visual targets. These results are consistent with the signals encoded by PPC neurons that I described in other chapters, and suggest that the PPC is necessary for the correct programming of more complex types of gaze shifts than those under direct visual guidance.

I

Modelling the Transformation of Spatial Coordinates

Chapter 2

Neural network models of coordinate transformations for gaze coding in the posterior parietal cortex

SUMMARY

In this chapter we review the use of a group of neural network models as tools for studying the function of a cerebral cortical area. The basic model is a feedforward multilayer network that learns to transform the coordinates of a visual stimulus from a retinocentric to a craniocentric reference frame using the backpropagation learning algorithm. The similarity of certain features acquired by the model's components with the response properties of neurons in the posterior parietal cortex made the model a candidate for studying the cortical area's processing in an artificial system. An extension of the model to one that transformed retinal coordinates into body-centered ones predicted response properties that were later confirmed by neurophysiological experiments. Simulation of electrical stimulation of the model predicted a pattern of effects similar to the one later obtained by stimulation of a specific region of the parietal cortex. More importantly, a study of the response properties of the model's units provided a simple explanation of how areas in the posterior parietal cortex might compute coordinate transformations. This analysis also suggested why certain manipulations such as stimulation should produce the effects observed. The same algorithm for coordinate transformation was also obtained in an analogous network trained with a learning rule biologically more plausible than backpropagation. These results suggest that neural network modeling is a useful adjunct to the neurophysiological and psychophysical techniques we are using to study the function of the posterior parietal cortex.

I. INTRODUCTION[1]

Two issues of active debate in systems neuroscience are what coordinate frames the nervous system uses to represent spatial information related to sensation and movement, and how it transforms spatial information among different coordinate frames. The nervous system of primates (and indeed of animals across the phyla) must represent spatial information in a variety of coordinates frames. This requirement is imposed at least by the fact that sensory organs and various body parts under motor control can move relative to one another. Successful spatially-oriented behavior demonstrates that the nervous system can transform signals among various coordinate frames. For example, when a person drinks coffee while reading the newspaper (Fig. 1a), she can reach for the coffee cup either by first shifting her gaze from the paper to look directly at the cup and then reaching for it (Fig. 1b), or—if the article is particularly absorbing—reach for the coffee cup without lifting her eye off the page (Fig. 1c). In the first case the gaze shift places the image of the cup directly on the fovea (labelled x in Fig. 1b), wile in the second case the cup's image falls on a different spot (labelled y in Fig. 1c) on the retina. The brain can program the correct reaching movement in either case, suggesting that sensory signals encoding the location of the cup on the retina are transformed into a reference frame appropriate for programming the reaching movement.

Much evidence suggests that a portion of the posterior parietal cortex (PPC) of the primate brain participates in the transformations of neural signals from sensory to motor coordinates. Specific areas within the PPC appear to compute such transformations for the purpose of programming the direction of gaze. We review here a set of neural network models developed in our laboratory to study how some of these coordinate transformations might be achieved in the PPC.

1. The material in this chapter has appeared in published form (Mazzoni and Andersen 1993; in press).

One of many uses of neural networks has been as models of neuronal ensembles that might give us some insight into the processing that neurons perform as a group. The approach in such studies, as in most modeling studies, is to construct a model that captures the features of a complex system that we believe are relevant to a particular function. One can then analyze, manipulate and modify the model with the hope of (1) gaining a further understanding of the system, (2) evaluating the role of various features in the system's function, (3) perhaps predicting other important features that could be observed experimentally, and (4) predicting the system's response to various perturbations. The network model that does incorporate some function or important feature of the neural system can then become a tool for expressing in explicit form various hypotheses and mechanisms related to the neural system.

The group of neural network models reviewed in this chapter have played the role just described in the study of the posterior parietal cortex (PPC) of the primate brain. The basic model is a feedforward multilayer network that learns to transform the coordinates of a visual stimulus from a reference frame anchored to the retina (retinocentric) to one centered on the head (craniocentric) (Zipser and Andersen, 1988). The network learns to vectorially add the stimulus' position on the retina to the position of the eye in the orbit. We will describe this model and its initial contribution to our understanding of the response properties of parietal neurons. We will then review the predictions produced by modifications and perturbations of the model, as well as the results obtained by similar manipulations and perturbations of the experimental system. The neural network model has so far helped not only to predict the system's behavior in various circumstances, but also to understand how its observed properties can subserve the computational functions attributed to this area of the brain.

II. Neuronal Properties And Presumed Function Of The Primate's Posterior Parietal Cortex

The PPC is thought to play an important role in the integration of sensory perception and motor behavior. Lesions of this area cause impairments of spatial abilities (reviewed in Andersen, 1987) and its neurons are specifically active during sensory stimulation and particular motor acts (reviewed in Andersen, 1987; Andersen and Gnadt, 1989). One requirement of sensorimotor integration is the transformation of spatial locations across coordinate frames. In order to reach for an object, for example, the location of its image on the retina must be transformed into the coordinate frame in which hand movement commands are generated (e.g. trunk-centered, shoulder-centered, etc.). These transformations are important for accurate spatial behavior because our sensory and motor organs can move relative to each other. Although our eye movements continually shift the image of the visual field on the retina we can still perceive a stable environment and make appropriate movements within it.

Early studies of the monkey's PPC (Fig. 2) revealed a group of neurons that responded to visual stimuli in an eye-position-dependent manner (reviewed in Andersen, 1987). The portions of the visual field in which luminous stimuli elicited responses, i.e. their receptive fields, corresponded to particular retinal locations. As the trained monkey rotated its eye to different gaze directions the receptive fields maintained their shape and retinal location, but the neurons' responses were modulated by eye position. Such eye-position-modulated receptive fields were referred to as "spatial gain fields," because eye position acted as a gain of the visual response (Fig. 3). A striking feature of these gain fields was that for a majority of the neurons the modulation was planar, i.e. proportional to the horizontal and/or vertical component of eye position, or had a planar component (Fig. 3c). The gain fields set PPC neurons apart from those of more peripheral visual areas, which encode strictly the retinal location of visual stimuli. However, PPC neurons do not code spatial locations

unambiguously at a single cell level; rather, the code for a stimulus' spatial location is distributed across the population response. If individual neurons were actually invariant for locations in head-centered coordinates, for example, then their receptive fields should remain fixed to some location relative to the head, i.e. their responses should be independent of eye movements. Instead, the individual neurons' firing rate is an ambiguous signal for stimulus location, because a change in activity can be due to a movement of the stimulus as well as caused by an eye movement. Another special feature of these neurons is that their receptive fields are extremely large, covering as much as half the visual field. Such a large response area makes it difficult for the individual neuron to code stimulus location precisely.

III. The Zipser-Andersen Model

The properties of PPC neurons suggested that individual neurons were unlikely to subserve the spatial integration functions attributed to this cortical region. Being sensitive to both retinal location and eye position these neurons remained candidates for playing a role in computing spatial relationships, but a spatial code could only be obtained from the pooled activity of a group of these neurons. Zipser and Andersen developed a neural network to study how an ensemble of neuron-like model units might solve the coordinate transformation problem (Zipser and Andersen, 1988). The aim was to examine the properties of individual units that were trained to solve the problem as a group. If the brain was indeed encoding spatial locations in the distributed pattern of activity of many parietal neurons, then some features of the brain's algorithm might emerge in the model network too.

The Zipser-Andersen model (Fig. 4) was a three-layer feedforward network of units with sigmoid input-output functions. In the input layer a group of units encoded the retinal location of a punctate visual stimulus and another group encoded eye position. Retinal location was

encoded topographically, each unit having a two-dimensional receptive field of Gaussian profile centered at a given point in the receptive field. Eye position was encoded linearly in the activity of units in the second group, using a different slope and intercept for each unit. These inputs were modelled according to input signals actually available to neurons in the PPC. Units in the hidden layer received signals from all the input units and projected to all output units. The output layer was intended to code for the head-centered location of the visual stimulus. The task was thus to perform the vector addition of two positions, the stimulus' retinal position and the eye position, to obtain the stimulus' location in a craniocentric reference frame. The output layer was trained to express head-centered locations in one of two formats: a receptive-field based one like the retinal input (topographic format) or a linear-function one like eye position input (monotonic format). The network was trained to compute the coordinate transformation from a set of examples using backpropagation. After the network learned the task, its hidden units were found to respond to visual stimuli and to eye position very much like PPC neurons. Specifically, they had retinotopic visual receptive fields whose activity profiles were modulated by eye position, i.e. they had spatial gain fields, and these gain fields were largely planar (Fig. 5). The receptive fields were also very large and smooth with one or a few eccentric peaks, and thus looked remarkably like those of PPC neurons (Fig. 6).

The results of this simulation addressed several issues related to PPC function. First, it was shown that a layered network can learn to transform retinocentric coordinates into craniocentric ones using the input signals available to the PPC (whether the training signal required by backpropagation is available to the PPC is unknown, but see the discussion of the reinforcement-based model below). This result is consistent with the adaptability of spatial behavior (e.g., when the visual input is distorted by prisms) that persists throughout life. In fact, such adaptation is easily elicited precisely by training subjects with examples, as in the change in amplitude of eye movements that is obtained when stimuli presented through distorting prisms. Second,

the representation of the spatial signals that emerged at the hidden units' level as a results of learning to solve this problem is very similar to the representation of the same signals in PPC neurons. These neurons can thus play a similar role in the organism, i.e. build up an intermediate representation between input and output stages that is part of the coordinate transformation computation. The network demonstrated explicitly that units with PPC neurons' properties contain, as a group, a distributed representation of space that is sufficient for accurate localization. Such a distributed representation obviates the need for a topographic representation of head-centered space. The version of the model trained with the monotonic output format demonstrated that the distributed spatial representation of the hidden layer can appropriately drive a non-topographic code directly. This output layer could represent a set of motoneurons driving the eye muscles to orient the eye toward the desired spatial location. The adequacy of the hidden layer's representation to feed directly into a peripheral output could explain why such a topographic map of head-centered or body-centered visual space is generally not found in the cerebral cortex. Finally, the fact that training a neural network to perform coordinate transformations produced a hidden layer with properties like those of PPC neurons suggested that the network and the brain may employ a common strategy in solving the problem.

IV. An Extension Of The Model: Transforming Retinal Locations Into Body-Centered Coordinates

Having the problem's solution programmed in a network model made it possible to further investigate what algorithms the PPC may indeed be using through analysis and manipulations of the model. An immediate question was how locations could be coded in other coordinate frames. The transformation from retinal to head-centered coordinates has a natural application in the programming of eye movements, the eyes having to move to particular positions relative to the head (though the actual coordinate frame in which eye movements

are programmed is still a matter of debate). There is in fact a subdivision of the PPC, the lateral intraparietal area (area LIP), that is directly connected to eye movement centers and that contains neurons active during the preparation of saccadic eye movements (Lynch et al., 1985; Andersen et al., 1990a, 1992; Blatt et al., 1990). These neurons have planar gain fields that allow them to encode, as a population, the head-centered location of a saccade target. Large gaze shifts, however, are achieved by coupled movements of the eyes and head; in this case the target's position must be calculated in body-centered coordinates. Evidence from lesion studies suggests that the PPC is necessary for the proper execution of not only eye movements but other forms of spatial behavior as well. Could the Zipser-Andersen network be modified to compute body-centered coordinates, and if so, what predictions would it make about the PPC?

Goodman and Andersen (1990) added a group of units encoding head position to the input layer of the Zipser-Andersen network and trained this new network to produce body-centered locations at the output layer. The new units encoded head position using linear functions of the orientation of the head along various axes, roughly simulating the signals that neck muscles' proprioceptors might produce. The network was otherwise equivalent to the original one. It learned to compute the correct body-centered location given retinal, eye, and head position inputs. The hidden units were found to be sensitive to all three input types. They had retinotopic receptive fields modulated by both eye and head position, each in a planar fashion. In other words, they developed planar "gaze fields," that is, linear modulation of visual responses along a particular direction of gaze, which is the sum of eye and head positions. Moreover, the "eye" gain field of a given hidden unit was always aligned (with the same direction and slope) with the same unit's "head" gain field. This was a natural solution for the network given the constraints of its architecture (the eye and head position inputs produced signals in very similar formats) and of the problem (eye and head position are indeed coupled for a given spatial position). The result suggested, however,

that if the PPC subserves coordinate transformations beyond the head-centered reference frame and does so with an algorithm analogous to the neural network's strategy, then it should contain units with gaze fields similar to those of the network just described. Such units have recently been identified in the PPC (Brotchie et al., 1991). Brotchie and Andersen trained monkeys to look in various directions by moving their eyes alone or by moving both their eyes and their head. A population of PPC neurons had visual responses modulated equivalently by eye or head position. These gaze fields were largely planar and the direction of eye and head position modulation was the same. This population was distinct from another group of PPC neurons whose visual responses were modulated by eye position alone.

The locations of visual stimuli are thus presumably encoded in body-centered coordinates by the group of neurons with gaze fields and in head-centered coordinates by the neurons sensitive with only eye position gain fields. Other neurons may exist in the PPC with other types of gain fields which may encode, for example, a stimulus' location in a reference frame appropriate for planning a reaching movement to that stimulus. Distinct neuronal populations in the PPC may thus encode the location of a sensory stimulus in various coordinate frames, allowing the animal to program a whole repertoire of movements relative to that stimulus.

V. How The Neural Network Transforms Coordinates

The addition of a head position input is only one example of the many manipulations that are possible once a neural network embodies a candidate of a biological solution to a problem. Various manipulations of this type led Goodman and Andersen (1990) to a simple explanation of how the network performs coordinate transformations. Over the course of learning each hidden unit develops a "preferred direction," that is, a direction in its input space along which to maximally modulate its activity. By maximal modulation we mean that an input vector parallel to the preferred

direction produces the unit's largest activation and a vector in the opposite direction produces the smallest activation (or largest inhibition from the resting activity level). This behavior is typical of hidden units in feedforward networks trained with examples. The activation of a typical hidden unit (linear, logistic, etc.) is a monotonic function of the cosine of the angle between its weight vector and the current input vector. The preferred direction is the direction of the unit's weight vector.

The units of the Zipser-Andersen network, receiving both retinal and eye position inputs, align their sensitivity in retinal space and in eye position space, and develop an eye position response field that approximates a plane oriented along what becomes the unit's preferred direction (we will call this direction $\mathbf{a}_i$ for the i^{th} hidden unit in the network). A hidden unit effectively collapses the multidimensional signal of the retinal and eye position units into two two-dimensional vectors, one for retinal and one for eye position ($\mathbf{r}$ and $\mathbf{e}$, respectively). The goal is to add these two vectors to obtain the head-centered position vector, $\mathbf{h}$, where

$$\mathbf{r}+\mathbf{e}=\mathbf{h}$$

A hidden unit's activation is proportional to the dot product of its input vectors and its preferred direction:

$$\alpha_i \approx \mathbf{r}\cdot\mathbf{a}_i+\mathbf{e}\cdot\mathbf{a}_i$$

Each hidden unit thus extracts the components of the retinal and eye position vectors along its preferred direction and adds them. Because these vectors' components are added at the hidden unit's input, the output of each hidden unit effectively consists of the component of the head-centered vector along the unit's preferred direction. Formally, because

$$\mathbf{r}\cdot\mathbf{a}_i+\mathbf{e}\cdot\mathbf{a}_i=\mathbf{h}\cdot\mathbf{a}_i$$

then

$$\alpha_i \approx \mathbf{h} \cdot \mathbf{a}_i$$

The preferred directions of the hidden units span the two-dimensional input space so that the retinal and eye position vectors are decomposed without losing information. These components are combined again at the output layer to give the head-centered vector that is the sum of the retinal and eye position vectors. A single hidden unit's operation can thus be described as a sum of dot products, and is an elegant way of adding two vectors that are encoded in the activity of many input units.

Understanding the algorithm discovered by the network model in these terms gives us a clearer view of how the brain might be solving the same problem. In this regard it is also useful to consider alternative solutions not employed by the network. One solution that the network does not choose is a "look-up table" hidden layer. In such a solution each hidden unit would link one or a small group of input vector pairs with their correct outputs in a dedicated (i.e. non-distributed) fashion. The network learns instead to perform the abstract computation of vector addition, as evidenced by its correct performance on novel pairs of retinal and position inputs. As the number of hidden units in the network is increased, a look-up table solution does indeed emerge. This observation suggests that a constraint on computational efficiency may be at work in the biological system's choice of solution.

The network also avoids solutions of the "shifter circuit" type (Anderson and Van Essen, 1987), whereby a group of hidden units might set up retinal receptive fields that are gated at the output layer by the signals of another group of hidden units concerned with eye position alone. The eye position signals in the shifter circuit serve to align the receptive fields of the visual input units with the those of the appropriate output units. Such a strategy solves the coordinate transformation problem just as well as a distributed intermediate representation, but the Zipser-Andersen network produces the latter. In fact, it seems a very common result that networks trained to

minimize an error signal at the output layer develop distributed representations in their hidden units. The fact that the PPC also encodes head-centered coordinates in a distributed fashion suggests that the brain arrives at this solution due to constraints analogous to those imposed on the model network, e.g. training with an error signal a group of neurons that are initially highly interconnected. Moreover, had the network used a shifter circuit solution, we would have predicted a different result when the head-position input was added: we would have expected the emergence of a separate cluster of head-position units, and not the distributed coding of head signals over all the hidden units. The gaze fields of the hidden units in the head-position network are thus not a chance prediction of the model, but something to be expected from the way a trained layered neural network computes coordinate transformations. The observation of neurons with gaze fields in the PPC then supports the hypothesis that the neural network model embodies some of the important features of the nervous system's solution to the same problem.

The most important factor in determining the hidden layer representation, as in any neural network, is the problem that the network must solve. The Zipser-Andersen network must learn to approximate a function, the vector addition of retinal and eye positions. The network's hidden layer representation, a distributed code of the input vectors' combined components, is one of the best solutions to this problem that one can imagine for the given architecture. It allows a few hidden units to encode a large number of input-output pairs, and it naturally produces interpolation for novel inputs. A look-up table solution would be more suitable for a classification problem such as deciding, for example, whether a visual stimulus can be reached with the left hand or with the right hand. The "shifter-circuit" solution was designed to *discount* the eye position to facilitate the analysis of the visual scene in the presence of eye movements. Using eye position to actually compute head-centered coordinates may be a different enough problem to make this solution not ideal. One can suppose, therefore, that PPC neurons have a

distributed representation of head-centered positions, manifested as a set of eye position gain fields, because this is a good solution to the coordinate transformation problem.

A notable feature of the distributed representation of the Zipser-Andersen network is the absence of topography in the hidden layer. Maintenance of topographic relationships across processing stages can be an effective mechanism for processing spatial information. Several models of saccade generation, for example, use representations with well-defined spatial relationships in order to generate an appropriate saccadic command to look at a sensory stimulus (e.g. Droulez and Berthoz, 1991; Dominey and Arbib, 1992). In these models the saccade vector is determined by *which* units are active within a given stage. Units in the hidden layer of the Zipser-Andersen network, on the other hand, are connected to every input and every output unit, and encode the head-centered position vector without regard to any input or output topography. The output vector is determined not by which units are active but by the activity level of every unit in the hidden layer.

It is not clear whether PPC areas are topographically organized. The Zipser-Andersen model demonstrates that PPC neurons can transmit to other cortical areas the head-centered position of a stimulus, encoded in their collective firing rate, without any topographic organization. Such organization, of course, may still be a feature of PPC without playing a role in the coordinate transformation computation. It may be a by-product, for example, of some developmental mechanism that establishes the initial connections of this region with other brain structures. In this regard an interesting modification of the Zipser-Andersen model was described by Cho and Reggia (1992). Their model was trained on the coordinate transformation problem with a learning scheme that combined backpropagation with competitive learning. The hidden layer acquired known features of PPC neurons as well as a topographic arrangement of craniocentric receptive fields. The topography was not necessary to

solve the coordinate transformation problem but it did not interfere with the solution either.

VI. Perturbing The Model: Simulating Electrical Microstimulation

If theoretical analysis of the neural network's behavior suggests that this model captures some features of the brain's algorithm for coordinate transformations, then it may be instructive to introduce perturbations to the model and observe its behavior. We have mentioned that one use of transforming retinal coordinates into head-centered ones is to compute the orientation of the eyes required to direct gaze to a peripheral visual target. One symptom of parietal lobe lesions is indeed difficulty in orienting gaze, and the PPC is directly connected with various centers controlling eye movements. Electrical stimulation of a cortical area is one way to test its role in the production of specific behavior. Simulating electrical stimulation in the neural network made specific predictions concerning the metrics of eye movements that would be obtained if the neural network's output encoded the programmed endpoint of an eye movement (Goodman and Andersen, 1989). These predictions were borne out by a detailed stimulation study in the monkey (Thier and Andersen, 1991).

The amplitude and direction of eye movements elicited from stimulation of a nervous system structure with the eyes starting at various orbital positions can elucidate how that structure represents spatial locations. Obtaining eye movements with metrics independent of the eye's initial position (fixed-vector pattern) suggests a coordinate frame centered on the eye, whereas amplitude and direction that vary with initial eye position so that the eyes are always moved to a single orbital position (convergent pattern) are consistent with a head-centered reference frame. Goodman and Andersen (1989) simulated the effect of electrical microstimulation by setting the output of a given hidden unit in a trained network to its maximum possible value and interpreting the new position encoded by the output layer as the

endpoint of the eye movement presumed to result from the stimulation. This process was repeated for many initial eye positions, each of which was fed to the network's input layer. What we know about the way the hidden layer encodes spatial locations allows us to predict the pattern of eye movements to be obtained from the network.

Because a hidden unit's activity encodes the component of the head-centered vector along the unit's preferred direction, maximal activation of that unit will shift the network's output along the unit's preferred direction. This direction is encoded in the unit's weights and so should not be affected by the values of the inputs; i.e., the direction of the eye movement will be independent of the starting eye position. The movement's amplitude, on the other hand, is proportional to the change in the unit's activation, which depends on how far from saturation the hidden unit is before stimulation. This follows from the units' sigmoidal input-output function. A hidden unit in the Zipser-Andersen network normally operates along the linear (center) portion of the sigmoid—which results in the planar component of the unit's gain field. Stimulation, however, drives the unit's output to its maximum value. The change in the unit's output therefore depends on how far the initial activation level is from the unit's maximum possible activation. The initial activation level is determined by the unit's input, and thus by the initial eye position. At starting eye positions whose vectors are orthogonal to the unit's preferred direction, the hidden unit will not be very active and maximal activation will produce a large shift in its output, giving a large-amplitude eye movement. At eye positions nearly parallel to the preferred direction the unit will be already very active and further activation due to stimulation will give only a small change in its output, and thus a small eye movement. (Note that this aspect of the unit's behavior is not inconsistent, as it may seem at first, with the increase in a neuron's responsiveness to visual stimuli as the eye position moves in the direction of the neuron's gain field. This responsiveness cannot be compared with the saturating effect of microstimulation because the background firing rate also increases as the eye position changes.)

The pattern of eye movements just described was indeed obtained by stimulation of most hidden units in a trained network (Goodman and Andersen, 1989). The eye movements had very similar directions from all starting eye positions, but their amplitude decreased as the eye position was shifted along one direction. The direction of this amplitude decrease was very similar to the direction of the elicited eye movement, indicating that the saccades were getting smaller as the eye moved along the unit's preferred direction, as predicted. Another prediction made in this study was that stimulation of two sites in the PPC should produce a much more convergent pattern of eye movements than stimulation of either site alone. This pattern is what we expect from the vector sum of the saccades elicited from each site.

In a detailed stimulation study of the PPC, Thier and Andersen (1991) found a correspondence between the types of eye movement that could be elicited and anatomical subdivisions of the PPC. Stimulation of area LIP (the region that directly projects to eye movement centers and that is active during the programming of saccadic eye movements) elicited a pattern of saccadic eye movements like those obtained from the neural network. The saccades evoked from various initial eye positions were all in the same direction, but their amplitude decreased as the starting eye position was moved in the direction of the elicited eye movement.

Neurons in area LIP also have spatial gain fields (Andersen et al., 1990b) and can thus compute head-centered locations to drive eye movements much as the neural network does. This hypothesis is supported by the pattern of eye movements obtained from simultaneous stimulation of two separate sites: the saccades evoked from various eye positions in this case converge toward a much more circumscribed region than in the case of single-site stimulation (Thier and Andersen, personal communication). The neural network model thus correctly predicted the pattern of eye movements to be expected from stimulation of a PPC region that had been independently implicated in the control of eye movements. The stimulation studies would have been performed in the monkey whether or not any neural

network model had existed or had made any predictions, as this technique has always been a classical neurophysiological tool for the study of nervous system structures. The particular effect predicted by double-site stimulation of the network at different starting eye positions did, however, prompt the experimenters to perform in the monkey this same experiment, which would not have been considered necessarily interesting otherwise. Moreover, the network provided an immediate and detailed explanation for why the eye movement pattern obtained—which suggests neither a retinal nor a head-centered representation—should indeed be expected from an area encoding spatial locations using gain fields.

Encoding the head-centered location of a stimulus is not the only way in which a saccade to that stimulus can be programmed. Another commonly proposed scheme maps the sensory vector falling on the retina (from the fovea to the stimulus' image) directly into a motor command encoding the required saccade vector, without ever computing the head-centered location of the stimulus. This method still requires some mechanism for keeping track of eye position, so that an appropriate saccade can be made to targets that appeared before one or more intervening eye movements. One such mechanism updates the planned saccade vector base on the last eye movement made. This method has been postulated as a cortical mechanism for saccade planning (Goldberg et al., 1990) and has been used in saccade-generation models (Dominey and Arbib, 1992; Droulez and Berthoz, 1991). In the models the future saccade vector is remapped based either on the integrated eye velocity signal from the intervening saccade (Droulez and Berthoz, 1991), or on a damped copy of the intervening saccade's eye position signal (Dominey and Arbib, 1992). The Zipser-Andersen model does not address the issue of multiple saccade plans. Extending the model to handle sequences of saccades, however, would not require a remapping scheme that kept track of intervening saccades. All saccade targets would be directly encoded in head-centered coordinates as they appear, and a saccade to each could

be planned based only on the current eye position, independently of past eye movements.

VII. Biological Plausibility Of The Learning Algorithm

The biological plausibility of the Zipser-Andersen model was an issue of concern because backpropagation, the learning algorithm used to train the model network, is an unlikely candidate as a biological learning mechanism. It had been argued that the hidden units of backpropagation networks acquire particular properties because this algorithm computes solution that are optimal in the sense of minimizing the output error (Zipser and Rumelhart, 1991. The hidden layer representations thus should not depend on details of the learning algorithm. Any algorithm that computes optimal solutions for an architecture analogous to the one of the biological network should produce hidden layer representations similar to the ones used by the brain.

To address this issue directly Mazzoni et al. (1991) trained a neural network to perform the retinal to head-centered coordinate transformation using a reinforcement learning rule developed by Barto and Jordan (the associative reward-penalty, or $A_{R\text{-}P}$, learning rule; Barto and Jordan, 1987). This study is described extensively in the next chapter. Briefly, the $A_{R\text{-}P}$ algorithm adjusts the network's connections based on (1) a single error value computed from the network's overall performance, and (2) the local presynaptic and postsynaptic activation for each connection. It combines, in other words, a reinforcement signal with Hebbian updating of connection strength, and is thus biologically more plausible than backpropagation. The hidden units of this network developed gain fields and receptive fields virtually identical to those of the backpropagation-trained networks. The networks' algorithm for computing coordinate transformations, therefore, did not depend on the specific learning mechanism used.

The $A_{R\text{-}P}$ learning rule, like backpropagation, implements a gradient-descent procedure that leads to the optimal solution in terms of minimizing the output error. The model of Mazzoni et al. thus demonstrated that the principle underlying the learning procedure is a more important determinant of the hidden layer representation than the learning rule itself. The fact that this model learned to compute coordinate transformations based on a simple reinforcement signal supported the idea that PPC neurons can learn to compute coordinate transformations from signals directly available to the nervous system. More importantly, the reinforcement-based model showed the coordinate-transformation algorithm obtained in the Zipser-Andersen model is not a specific result of backpropagation learning. Thus the use of backpropagation in the networks we have described does not invalidate our interpretation of these networks as models of PPC function.

VIII. CONCLUSIONS

The network model of PPC developed by Zipser and Andersen has proved a valuable tool in our study of this cortical area's function. Like every useful model it embodies a few known features of the system being studied and it is helping us to understand why those features, and others that emerge from the model, are important for the system's computations. The model has predicted a few experimental results and has made it easy to put those results into an explicit theoretical context. The general approach has consisted of a combination of a novel modeling paradigm and a "reverse engineering" strategy of brain function investigation. A fundamental tenet of neurophysiology is that we can try to understand the functioning of the brain by observing the responses of its components to various stimuli and under various conditions. This method is a form of reverse engineering much like studying how a typewriter works by watching its parts move when someone presses the keys. The response properties of PPC neurons are complicated enough that reverse engineering applied directly to the

brain was not sufficient, and a model was thus devised. A neural network was chosen as a model because it can program itself for an optimal solution given a problem (coordinate transformation) and a set of input and desired output signals (the signals that feed into the PPC and accurate spatial behavior, respectively). Reverse engineering was then applied to the trained neural network, with the advantage that the researchers could perform the manipulations necessary to understand the algorithm arrived at by the network. The properties and behavior of the biological network could thus be explained within a theoretical description of its processing function.

We believe the neural network has been a useful modeling paradigm for studies of the PPC especially because these studies addressed how neurons encode more than one parameter. Whereas most neurophysiological investigations so far have focused on the effect of varying one parameter at a time (a necessary first step for the study of any complex system), studies of an area thought to compute spatial relationships had to examine the interactions of several variables such as retinal location, eye position, and head position. It is not surprising then that the experimental data was not easily summarized by an intuitive scheme. The network model provided a framework for developing an intuition about the distributed representation of several variables. As more experiments address the encoding and interactions of several parameters in the nervous system, we expect neural networks to continue to fruitfully assist our investigations of nervous system functions.

FIGURE LEGENDS

Figure 1:

Example of a behavior that requires the transformation of spatial sensory signals. A person reading a newspaper (*a*) can reach for a coffee mug either by first looking at it (*b*) or without detaching gaze from the page (*c*). In each instance the mug's image falls on a different location on the retina (at x in *a* and at y in *b*), but the required reaching movement is the same.

Figure 2:

a) Lateral view of the right hemisphere of a macaque monkey's brain (*Macaca mulatta*). Anterior is to the right. The intraparietal (IP) and superior temporal (ST) sulci are labelled. The thick solid line marks the approximate boundaries of the PPC. The shaded region is area 7a, one of the visual areas of the PPC. Area LIP (not visible) is buried in the intraparietal sulcus, adjacent to area 7a. *b*) Visual pathways from the striate cortex (area V1) into the PPC (areas enclosed by dashed rectangle) in the macaque monkey. These areas form the bulk of the "where" pathway (i. e. the pathway that analyses spatial visual information) of the primate visual system. (Adapted from Andersen, 1987)

Figure 3:

Measurement of a neuron's gain field. *a*) A monkey trained to maintain gaze on a fixation light spot sits in a dark room facing a projection screen while the activity of single cortical neurons is recorded. While the monkey maintains fixation a second light spot is used as a probe to locate the region of the visual field where a light stimulus elicits a response from a particular neuron (its receptive fields, *rf*). *b*) The location, relative to the fixation point (and thus relative to the fovea), giving the peak response is then stimulated while the animal's angle of gaze is varied by placing the fixation spot at

different locations. Here the neuron's response to a stimulus presented at a constant retinal location varies as the eye moves among nine different fixation positions. Each histogram is a plot of the neuron's average firing rate vs. time, over the duration of stimulus presentation (500 ms), plus a brief period before and after. The relative positions of the histograms correspond to the nine different fixation positions tested. These positions were spaced by 20 deg., with the center one positioned straight ahead. *c*) Another way to represent the effect of eye position on the neural response is to sum the total response to the stimulus and represent it as the diameter of a filled circle. The circles, like the firing rate histograms, are arranged according to the eye position to which they correspond. (Adapted from Zipser and Andersen, 1988)

Figure 4:

Structure of the Zipser-Andersen neural network. A) Network architecture. The input layer (bottom) consists of 64 units encoding retinal position topographically and 32 units encoding eye position via linear functions. All input units have feedforward connections with all the hidden units, which have logistic input-output functions. Each hidden unit in turn projects to all output units. These are 32 logistic units (in the monotonic format; the topographic output format is not shown) that are trained to encode the vector sum of the retinal and eye position locations encoded in the input layer. B) One-dimensional section of the position-encoding function for a retinal input unit. Each unit's activity is a gaussian function of the retinal x and y coordinates of the visual stimulus. The gaussian functions represent visual receptive fields and are arranged in an 8x8 array. C) Linear position-encoding function for an eye position input unit. Each unit's function has a different (pseudo-randomly chosen) slope and intercept. D) Schematic of a logistic unit. It receives a set of inputs (the vector $\mathbf{x}$) which are multiplied by weight values (the unit's weight vector $\mathbf{w}$) and added together. The resulting value s (the dot product of $\mathbf{x}$ and $\mathbf{w}$),

together with a bias b, is passed through the sigmoid function $1/(1+\exp(-s))$ to yield the unit's output y. E) Linear position-encoding function for the output units. As for the eye position inputs, each unit in the output layer has a position function with a different slope and intercept.

Figure 5:

Gain fields of PPC neurons and hidden units in the Zipser-Andersen neural network. The format is the same as in Fig. 3c. *a*) Total response to a visual stimulus with the eye at each of nine position is shown for four neurons in area 7a of the PPC. *b*) Activity of four hidden units of a trained Zipser-Andersen network during the presentation of a visual stimulus with the eye at each of nine eye position. (Adapted from Zipser and Andersen, 1988)

Figure 6:

Visual receptive fields of PPC neurons and hidden units in the Zipser-Andersen neural network. Each field shows the response to a visual stimulus presented at several locations within a 40 deg. radius circle around the fixation point. *a*) Receptive fields of five area 7a neurons. These response surfaces were interpolated from 17 sampled points (see Zipser and Andersen, 1988). *b*) Receptive fields of five hidden units of a trained Zipser-Andersen network. These were sampled at 10 deg. intervals within the 40 deg. radius circle. (Adapted from Zipser and Andersen, 1988)

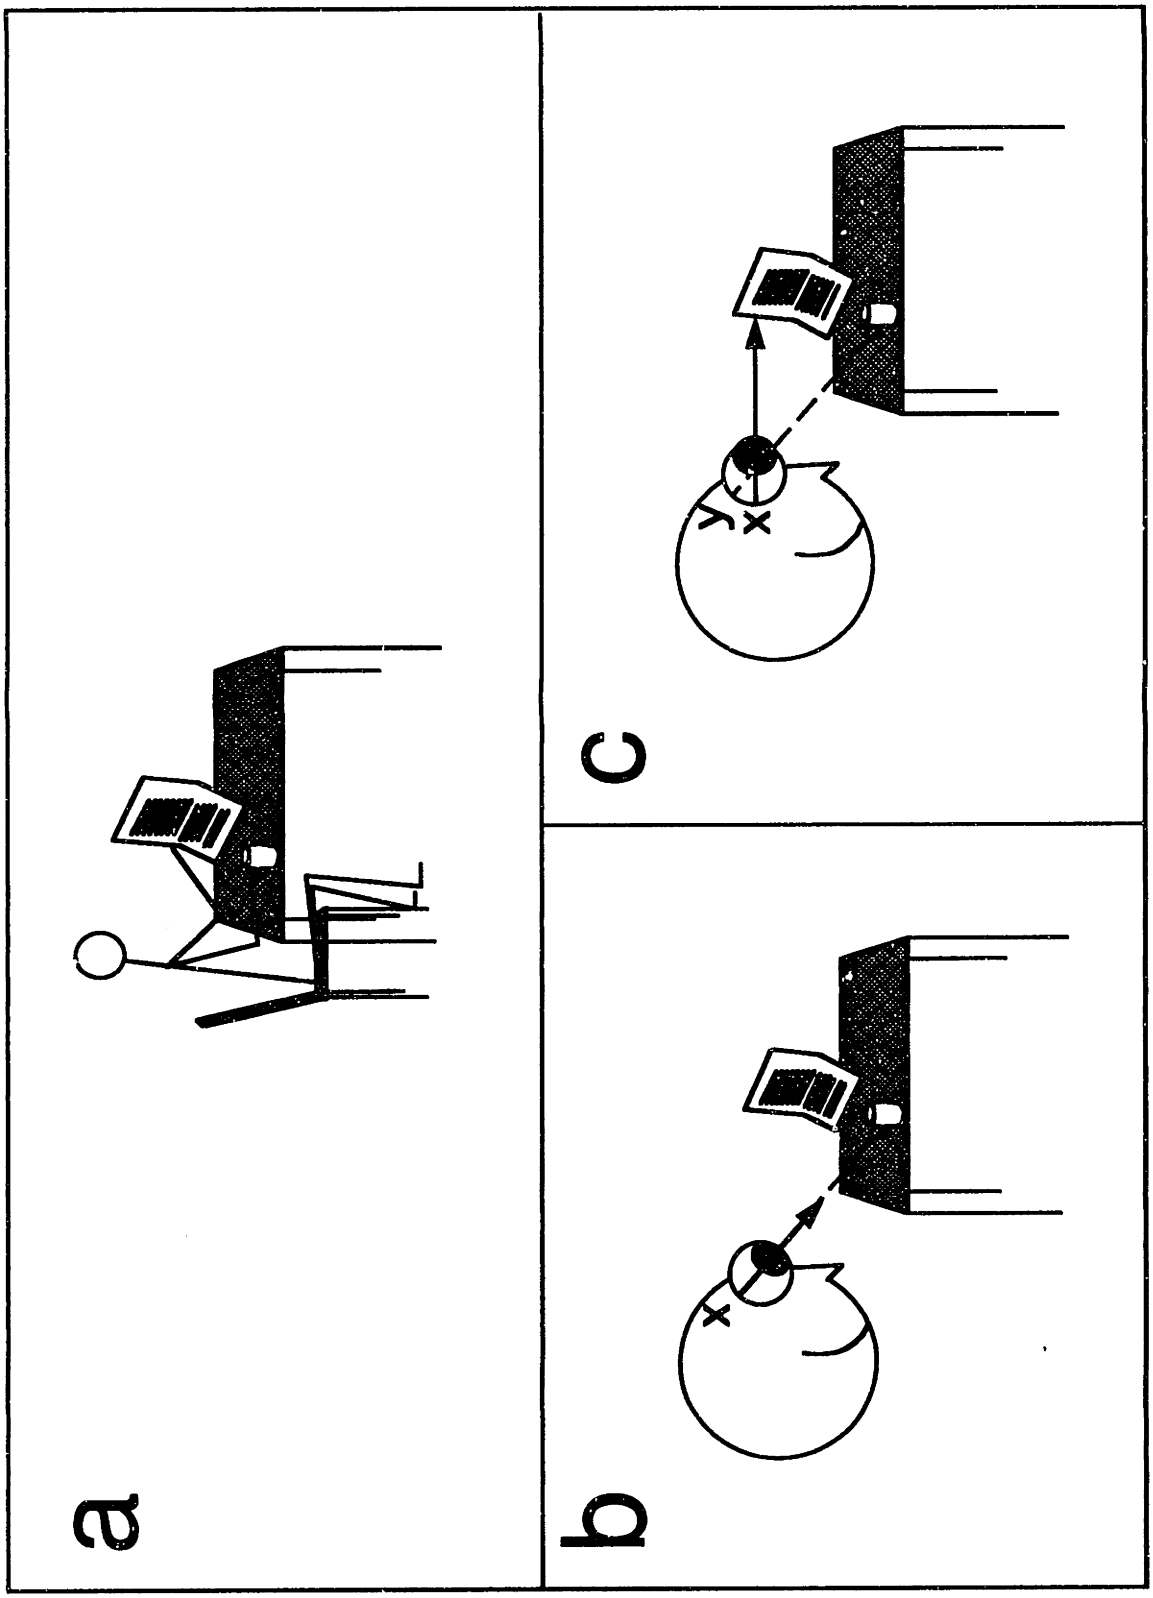

Figure 1

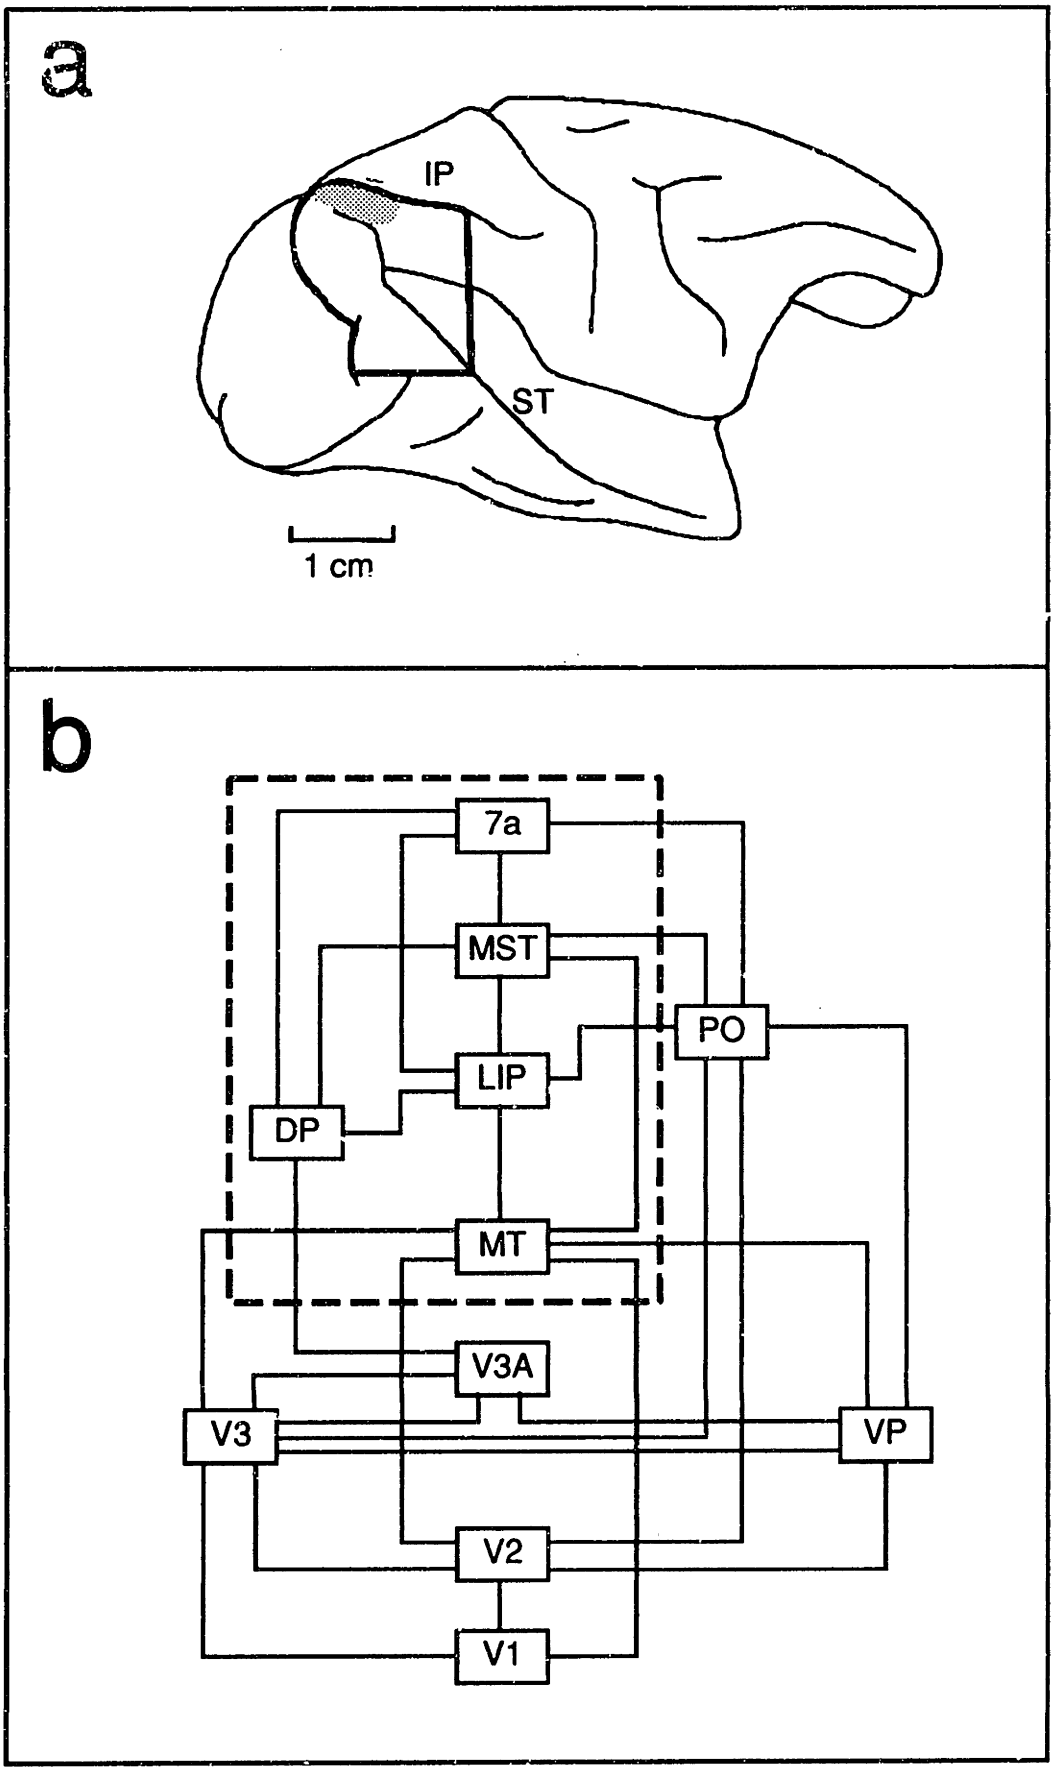

Figure 2

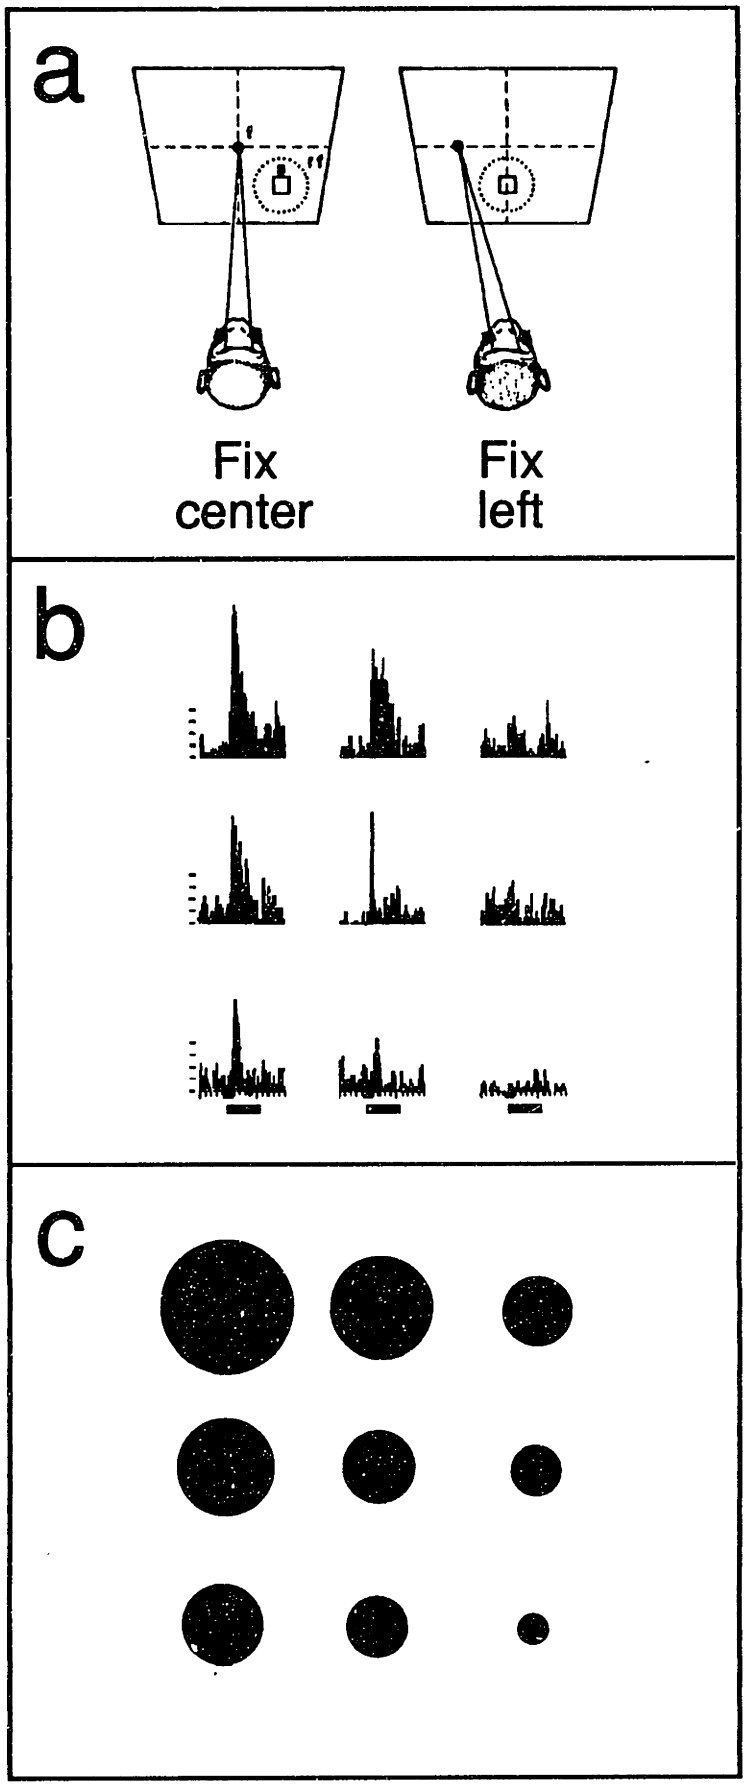

Figure 3

a
Output = Head-centered Location
head x (+ slope)
head x (- slope)
head y (+ slope)
head y (- slope)
W
Hidden Units
W
W
Eye Position
eye x (+ slope)
eye x (- slope)
eye y (+ slope)
eye y (- slope)
retina y
retina x
Retinal Location
b
activity
retinal position
c
eye position
d
b
w_o
w_j
w_n
x_o
x_j
x_n
Σ
s
y
e
head-centered position

Figure 4

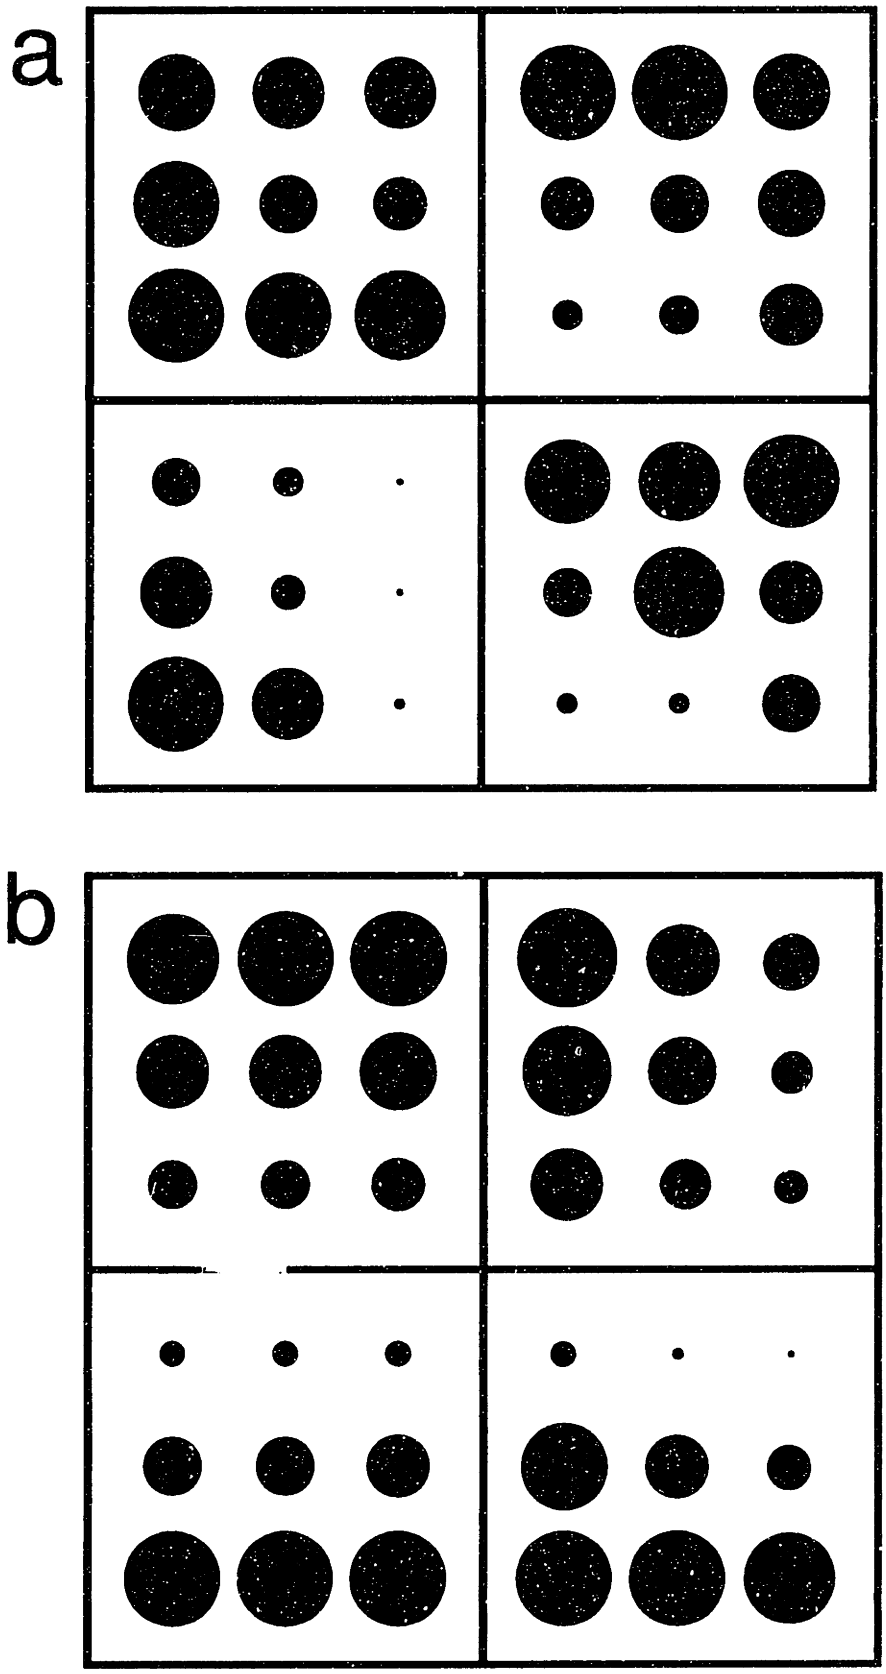

Figure 5

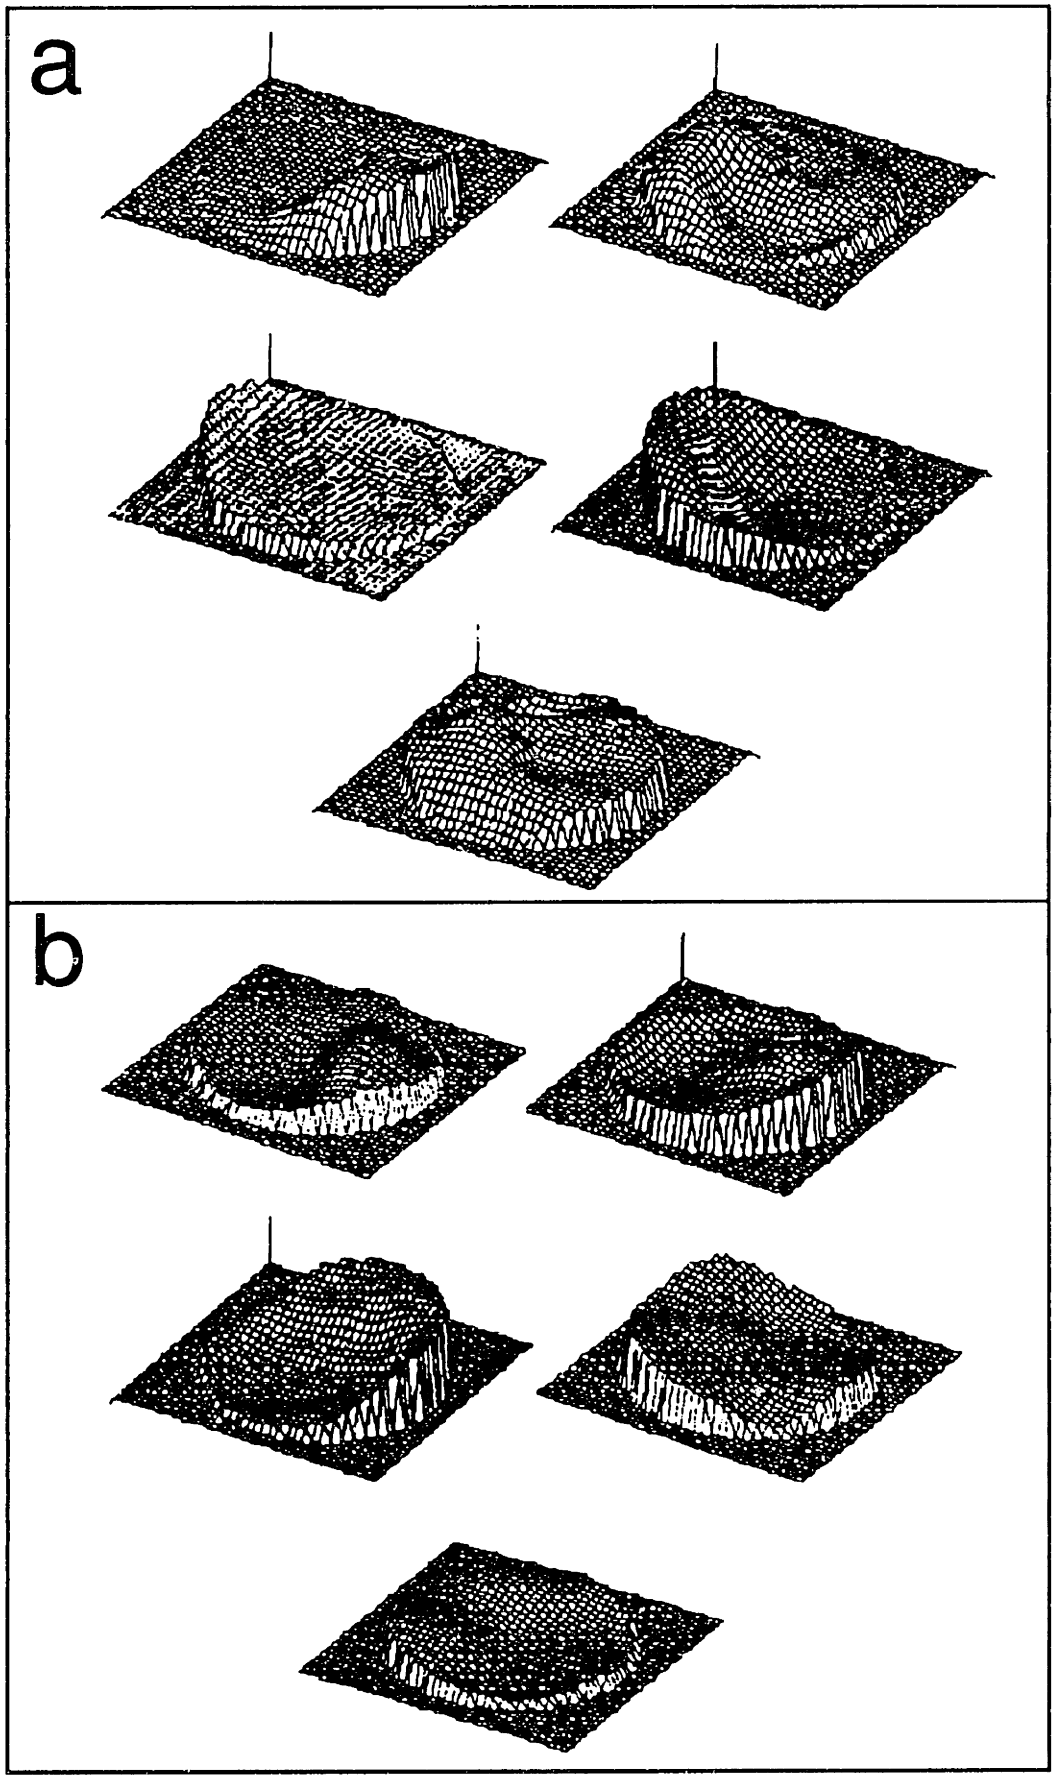

Figure 6

Chapter 3

A more biologically plausible model of coordinate transformations in the posterior parietal cortex

SUMMARY

Area 7a of the posterior parietal cortex of the primate brain is concerned with representing head-centered space by combining information about the retinal location of a visual stimulus and the position of the eyes in the orbits. An artificial neural network was previously trained to perform this coordinate transformation task using the backpropagation learning procedure, and units in its middle layer (the hidden units) developed properties very similar to those of area 7a neurons presumed to code for spatial location (Andersen and Zipser, 1988; Zipser and Andersen, 1988). We developed two neural networks with architecture similar to Zipser and Andersen's model and trained them to perform the same task using a more biologically plausible learning procedure than backpropagation. This procedure is a modification of the *Associative Reward-Penalty* ($A_{R\text{-}P}$) algorithm (Barto and Anandan, 1985), which adjusts connection strengths using a global reinforcement signal and local synaptic information. Our networks learn to perform the task successfully to any degree of accuracy and almost as quickly as with backpropagation, and the hidden units develop response properties very similar to those of area 7a neurons. In particular, the probability of firing of the hidden units in our networks varies with eye position in a roughly planar fashion, and their visual receptive fields are large and have complex surfaces. The synaptic strengths computed by the $A_{R\text{-}P}$ algorithm are equivalent to and interchangeable with those computed by backpropagation. Our networks also perform the correct transformation on pairs of eye and retinal positions never encountered before. All of these findings are unaffected by the interposition of an extra layer of units between the hidden and output layers. These results show that the response properties of the hidden units of a layered network trained to perform coordinate transformations, and their similarity with those of area 7a neurons, are not a specific result of backpropagation training. The fact that they can be obtained by a more biologically plausible learning rule corroborates the validity of this neural network's computational

algorithm as a plausible model of how area 7a may perform coordinate transformations.

INTRODUCTION[1]

An important element of information processing in the nervous system appears to be the collective behavior of large ensembles of neurons. The study of the emergent properties of these networks has been an important motivation behind the development of artificial neural network models whose architecture is inspired by the biological wiring of nervous systems, containing a large number of simple computational units extensively connected to one another. It is the hope of many neuroscientists that these models will elucidate, at least at an abstract level, some of the basic principles involved in information handling by the nervous system, and thus perhaps provide a theoretical framework within which to formulate experimental questions.

One of the best examples of this type of approach so far is a neural network model of area 7a of the primate's posterior parietal cortex developed by Zipser and Andersen (1988; Andersen and Zipser, 1988). From lesion and single-cell recording studies in primates it appears that area 7a is concerned with the representation of spatial locations in a head-centered reference frame (see Andersen, 1989, for a review). This representation is distributed over a group of neurons which are sensitive to both the position of the eyes in the orbits and the location of visual stimuli on the retinas. Other neurons in area 7a respond to either eye position or visual stimuli alone, and are presumed to provide the inputs from which the visual/eye position neurons extract the craniotopic representation. The latter neurons have very large retinotopic visual receptive fields and their response to eye position interacts nonlinearly with the visual signals. Although the majority of area 7a neurons maintain the same retinotopic receptive fields for different eye positions, the magnitude of the visual response varies with angle of gaze. Holding the retinal location of a visual stimulus constant and varying eye position (Fig. 1a), Andersen and his

1. The material in this chapter has appeared in published form (Mazzoni et al. 1989, 1991a, 1991b).

colleagues found that these neurons' overall firing rate (visual plus eye position components) varied roughly linearly with changes in horizontal and vertical eye position (Fig. 1b). The response profiles for varying eye position were called "gain fields," and a majority (78 percent) of area 7a cells had planar or largely planar gain fields (Fig. 1c; Andersen et al., 1985b; Andersen and Zipser, 1988; Zipser and Andersen, 1988).

Zipser and Andersen (1988) designed a computer-simulated neural network with an input layer, a layer of internal or hidden units, and an output layer. The input layer consisted of two groups of units with properties similar to those of area 7a neurons sensitive to either eye position or visual stimulus alone. The output layer coded for head-centered location in an abstract format independent of eye position, and was used to generate error signals to train the network. The network was trained to perform the coordinate transformation from retinotopic to craniotopic reference frames using the backpropagation procedure (Rumelhart et al., 1986a). The striking result of these simulations was that in the process of learning the hidden units developed response properties very similar to those of the area 7a neurons that seem to encode spatial location—specifically, a roughly planar modulation of visual response by eye position, and large complex receptive fields. This result suggested that area 7a neurons, as an ensemble, can in fact provide information for the abstract representation of space, and that these neurons' properties can be generated by a supervised learning paradigm.

Backpropagation networks can learn to perform a computation without explicit "knowledge," using only error signals from the environment as cues to improve its performance, in a paradigm referred to as "supervised learning" (Hinton, 1987). This type of training scheme has conferred upon neural networks a stronger element of biological plausibility than many previous models of brain function that relied on precompiled rules and symbol processing. Moreover, although various supervised learning algorithms for one-layer networks were described long before backpropagation (Minsky

and Papert, 1969; Widrow and Hoff, 1960), the discovery of the backpropagation algorithm (Werbos, 1974; Parker, 1985; Rumelhart et al., 1986a), which can be applied to more powerful multilayer networks composed of nonlinear units, is in large part responsible for the recent increase in interest in neural network models. In spite of the biological plausibility of supervised learning, however, the implementation of backpropagation in the nervous system would require mechanisms, such as the retrograde propagation of signals along axons and through synapses and precise error signal that are different for each neuron, which are not accepted as likely candidates for learning processes in the brain. To this end, Zipser and Andersen emphasized that it was the solution that was of interest in their model and not the method by which this solution was learned. They speculated that other, more biological learning procedures might generate a solution to the coordinate transformation task similar to that which resulted from backpropagation learning. It was therefore important to ask how crucial is backpropagation for the development of the hidden units' properties in a model like Zipser and Andersen's.

We addressed precisely this question in our study. We trained two neural networks with architecture similar to the Zipser and Andersen model using a supervised learning paradigm that is more plausible from a biological perspective than backpropagation. The algorithm we used, which is a variant of the *Associative Reward-Penalty* ($A_{R\text{-}P}$) algorithm for supervised learning tasks introduced by Barto and Jordan (1987), trains a neural network using a global reinforcement signal broadcast to all the connections in the network. We found that our networks can indeed be trained by these algorithms to perform the coordinate transformation task, and that the hidden units acquire response properties very similar to those of area 7a neurons, as in the Zipser and Andersen model. A second layer of hidden units can be interposed between the original hidden layer and the output layer without affecting the properties developed by units in the first hidden layer. Furthermore, all of our networks perform the correct

transformation on pairs of eye and retinal position never encountered before, that is, they generalize appropriately.

MATERIALS AND METHODS

We devised two types of networks which we trained to map visual targets to head-centered coordinates, given any arbitrary pair of eye and retinal positions. The basic architecture of these networks is similar to that of Zipser and Andersen's model.

Mixed A_{R-P} network

We call the first network the *Mixed A_{R-P} network* (Fig. 2a), because its hidden and output layers are trained by different learning rules (see Training, below). It is composed of three layers of computing units: an input, a hidden, and an output layer. The network has a fully connected feedforward architecture, meaning that every unit in each layer sends a signal to every unit in the next layer through an individual connection strength or weight (w), so that signals propagate in one direction from the input towards the output layer. The input layer consists of two groups of units (Fig. 2a, squares), one coding for the retinal location of the visual stimulus, and the other for the position of the eyes in the orbits. These units encode the external input by transforming an angular position into a value between zero and one, which is then sent to the hidden units. Retinotopic locations are represented by 64 visual units arranged in an 8x8 array, each with a gaussian receptive field (Fig. 2b) with peak at 10 degrees from its neighbors' and $1/e$ width of 15 degrees, producing a uniform topographic representation of the retina. Eye position is coded by 4 sets of 8 units representing horizontal and vertical eye coordinates with positive and negative slopes, for which activation is a linear function of eye angle (Fig. 2c). These representation formats were modeled according to characteristics of area 7a neurons established in previous studies (Andersen and Zipser, 1988; Zipser and Andersen,

1988; Andersen et al., 1985b), and are the same as those used in the input layer of the Zipser and Andersen model.

The hidden layer (Fig. 2a, diamonds) is so described because it is not "visible" (i.e., directly connected) to external agents acting at the input or at the output. The type of computational unit making up this layer is the *binary stochastic element* (Fig. 3a). This probabilistic element performs a weighted sum (s) of its inputs and passes it through the logistic function[1] to obtain a value (p) between zero and one. This value is then used as the probability of producing an output equal to one. The output is zero with probability $1\text{-}p$. This type of computing element is "neurally inspired", in the sense that it incorporates some well-established features of neurons. In such an analogy the inputs correspond to synaptic inputs from other neurons, the connection weights represent synaptic strengths (with inhibitory synapses implemented as negative weights), and the weighted input represents the intracellular potential. The probability of firing is analogous to a neuron's rate of action potential firing, and changes in the total weighted input affect the unit's probability of firing in a manner similar to how changes in the intracellular potential affect a neuron's firing rate. This hidden layer differs from that of the Zipser and Andersen network in that the units of the latter were deterministic logistic elements (described below). The number of hidden units in the Mixed $A_{R\text{-}P}$ network, as well as in all the networks described below, varied from two to eight.

The hidden units project in turn to the output layer, which encodes the craniotopic location that is the vector sum of the positions encoded by the retinal and eye position inputs. The units in the output layer (Fig. 2a, circles) are *deterministic logistic elements* (Fig. 3b). Like the binary stochastic units, they too perform a weighted sum of their inputs and pass it through the logistic function. In the deterministic

1. The logistic function, which is a type of "squashing" function, has a sigmoidal shape and maps real-valued inputs into the interval 0 to 1, according to: $f(s) = 1 / (1 + \exp(-s))$. In our networks s is the sum of the unit's inputs weighted by the corresponding connection strength, plus a bias.

logistic element, however, this value between zero and one is the unit's output itself. The output, therefore, is continuous and precisely predictable from the input. In the analogy with the neuron, this continuous output would correspond to the firing rate. The outputs of the output units encode head-centered location in one of two possible formats: a "monotonic" representation analogous to the eye position input, containing any number of units from 2 to 32 ("output 1", Fig. 2a), and a "gaussian" representation similar to that of the retinal input, with a number of units ranging from 4 to 64 ("output 2", Fig.2a). In the monotonic representation the activity of the output units increases for more peripheral locations of visual target with respect to the head, regardless of eye position. The gaussian format units fire for visual stimuli appearing within limited receptive fields in head-centered coordinates. We used either representation interchangeably, as this did not seem to affect our results. A physiological correlate of the monotonic representation could be a motor signal to the extrinsic eye muscles (Zipser and Andersen, 1988; Goodman and Andersen, 1989), while the gaussian format would be more like a receptive field for a mental representation of craniotopic space. These output representations are similar to those of the Zipser and Andersen model.

All-$A_{R\text{-}P}$ network

The second type of network we studied is the *All $A_{R\text{-}P}$ network* (Fig. 4a), so called because all of its connections are adjusted by the $A_{R\text{-}P}$ algorithm (see Training, below). This network's input and hidden layers are identical to those of the Mixed $A_{R\text{-}P}$ network. The output layer, however, is composed of binary stochastic units like the hidden layer. It too encodes craniotopic location in one of two alternative formats. Due to the binary nature of the output units, we devised output formats for the All $A_{R\text{-}P}$ network such that the collective activity of the output units codes for discrete adjacent regions of space, instead of continuously varying spatial locations. In the "binary-monotonic" format, 4 triplets of units divide craniotopic space into 16 regions by giving an output of one if the x (or y) head-centered coordinate is

greater (or less) than -40, 0, or +40 degrees (Fig. 4b). This format is analogous to the monotonic format of the Mixed $A_{R\text{-}P}$ network. The binary counterpart of the continuous gaussian output is the "binary-gaussian" format, in which 4 units have overlapping receptive fields centered at (±60, ±60) degrees, such that each unit's output is one if the spatial position is within 100 degrees of its center (Fig. 4c). This format divides craniotopic space into 13 regions by virtue of the overlap of the output units' receptive fields. The number of units in both types of binary output format may be increased to improve the output's spatial resolution. We did not examine the parameter of number of output units systematically, as it did not produce significantly different network behaviors.

Other networks

In addition to the two three-layer networks just described, we studied the behavior of two similar four-layer ne tworks. These consist of a Mixed $A_{R\text{-}P}$ network and an All $A_{R\text{-}P}$ network, each with an extra layer of hidden units between the first layer of hidden units and the output layer. This extra layer, like the original hidden layer, is also composed of binary stochastic units. We did this to see whether the response properties developed by the units in the hidden layer of the three-layer networks depended on a direct connection with the output layer.

For comparison purposes, we also set up a backpropagation network identical to the Mixed $A_{R\text{-}P}$ network described in Fig. 2a, except in its hidden units, which are deterministic logistic elements and not binary stochastic elements.

Training

We trained our networks to perform the coordinate transformation task in a supervised learning paradigm. In supervised learning, the network starts out with all connections and biases set at zero, or at some set of small random values if the training algorithm cannot break the initial symmetry (the $A_{R\text{-}P}$ algorithm we used can handle both

cases). An input pattern is presented to the network's input layer, which propagates a signal to the following layers (Fig. 5, solid arrows). The output layer thus produces a "guessed" output based on the initial set of connections. This output is compared to the correct output pattern for that particular input, and an error is computed and fed back to the network (Fig. 5, dashed arrows). The values of all the network's weights and biases are then adjusted by a learning rule so that at the next presentation of the same pattern the error is, at least on the average, smaller than before. This procedure is repeated until the error is reduced to a value below a desired level.

For our task the input pattern is a signal for the retinal location of a visual stimulus paired with one for the current eye position. The desired output pattern is one that codes for a head-centered location that is the vector sum of the retinal and eye positions. The error signal is computed externally to the network. To draw an analogy with how an animal may learn the coordinate transformation task, the input pattern would correspond to a visual stimulus seen with the eyes at a known angle of gaze (sensed by proprioceptive or corollary discharge pathways). The animal may then make a movement toward the stimulus, and any metric of performance, such as visually detected inaccuracies, could be used to generate an error signal.

The network is trained by being repeatedly presented with a finite number of patterns forming a chosen training set, the connection weights being adjusted after each pattern presentation. We used two types of pattern sets to train the networks. One is a set of random pairs of retinal locations and eye positions so that the desired output associated with each input is an arbitrary location in head-centered space. In the analogy with the learning animal, learning with this set would correspond to looking at various stimuli in the visual field at various angles of gaze. The other type of training set consists of input patterns for which the eye position is chosen randomly, while the retinal location is computed so that the vector sum of the two inputs is one of a few chosen craniotopic locations. The resulting training set contains a few fixed spatial locations, each represented by a large

number (at least ten) of retinal and eye positions that add vectorially to it. For an animal, this type of training corresponds to looking at an object fixed in space with the eyes in various orbital orientations. This training set was used to see how well the network generalized to new locations in space once it had been trained on a few fixed ones.

We devised two variants of the supervised learning procedure for $A_{R\text{-}P}$ networks of Barto and Jordan (1987) to adjust the weights of our networks. The essence of this algorithm is the $A_{R\text{-}P}$ learning rule. Every binary stochastic element in a given network receives a scalar *payoff* (or reinforcement) signal, r , whose value, in the supervised learning paradigm, depends on how close the current output is to the desired output. Specifically, r assumes a value between zero and one, with zero indicating maximum error in the output (i.e., every unit that should be firing is not, and vice versa), and one corresponding to optimal performance (no error in the computed head-centered position). The weights of the input connections on each binary stochastic element are then adjusted in such a way as to maximize the value of this payoff. Using the notation of Figure 3a, where x_i represents the output of the ith unit in the network, p_i its probability of firing, and w_{ij} the connection weight for its input from the j^{th} unit, the equation for updating the weights on a binary stochastic unit is

$$\Delta w_{ij} = \rho r(x_i - p_i)x_j + \lambda\rho(1-r)(1-x_i-p_i)x_j \qquad (1)$$

where Δw_{ij} denotes the change in the value of the connection strength w_{ij} after each pattern presentation, and ρ is a constant parameter representing the learning rate. As shown in Figure 3a, each unit also has a constant input, or bias (b_i). The value of this bias is also adjusted by the rule in Eq. 1, setting $x_j = 1$. Typical values for the parameters in this equation were 0.3 for ρ and 0.01 for λ. We will

describe this equation in more detail in the Discussion. The value of r is computed, externally to the network, as

$$r = 1 - \varepsilon \tag{2}$$

where ε is a measure of the current error at the output layer. In our model, ε is computed as the n^{th} root of the absolute value of the output units' error averaged over the number of output units:

$$\varepsilon = \left\{ \frac{1}{K} \sum_{k=1}^{K} \left| x_k^* - x_k \right| \right\}^{\frac{1}{n}} \tag{3}$$

where k indexes the K output units in the network, x_k^* is the desired output of the k^{th} unit in the output layer, x_k is its actual output, and n is a constant. Values for n ranged from 2 to 6. This expression for ε is different from the one used by Barto and Jordan (1987) , who computed ε as the sum of the squares of the output units' errors. Both expressions give a quantity nonlinearly related to the absolute value of the output units' errors, but our expression is more sensitive to small errors (a given unit's absolute error, $x_k^* - x_k$, is always less than or equal to one). Barto and Jordan referred to their learning algorithm as the "S-model $A_{R\text{-}P}$ rule", borrowing terminology from learning automata theory. In order to distinguish our modification of this rule from the original one we refer to our training algorithm as the *S'-model $A_{R\text{-}P}$ rule.*

We used the S'-model $A_{R\text{-}P}$ rule to adjust the weights of all the binary stochastic units in our networks. This includes all the weights in the All $A_{R\text{-}P}$ networks, and the weights between the input and hidden layers of the Mixed $A_{R\text{-}P}$ network. The output units of the Mixed $A_{R\text{-}P}$ network, being deterministic units with continuous output,

were trained by the delta rule for output units (Rumelhart et al., 1986a). This rule adjusts the weights of each output unit according to

$$\Delta w_{ij} = \alpha[(x_k^* - x_k) f'(s_k)] x_j \quad (4)$$

where k indicates the k^{th} output unit, α is a scalar learning rate, f' is the derivative of the logistic function with respect to the unit's net input s_k, and x_j is the output of one of the hidden units "presynaptic" to it. Typical values for α were between 0.5 and 2.5. This learning rule also is the basis for the backpropagation algorithm, and is used in identical form to train the output units of backpropagation networks.

RESULTS

Learning

All the networks described above learned to perform the coordinate transformation task to any desired accuracy. Figure 6 shows the general behavior during training of the two $A_{R\text{-}P}$ networks studied and compares them to that of a backpropagation network, with identical architecture, learning from the same training set. We plot here the absolute value of the output units' error, averaged over the number of output units and the number of patterns in the training set, versus the number of presentations of the complete training set. The $A_{R\text{-}P}$ networks produce learning curves with much more jitter than the curve for backpropagation training, due to the stochastic nature of their hidden units and to the type of error signal used in $A_{R\text{-}P}$ training (see Discussion). All three networks, however, produce curves with similar envelopes, and the times required for convergence are comparable. For the backpropagation network, which has a continuous output, the error decreases monotonically (Fig. 6a), while for the All $A_{R\text{-}P}$ network, which has a binary output, the error follows a noisy path down to zero and spends increasingly more time there, flickering occasionally to the value of the output's smallest resolvable angle (Fig.

6c). The error for the Mixed $A_{R\text{-}P}$ network is also noisy, because this network's hidden units are stochastic. It assumes, however, a continuous range of values (Fig. 6b), because the output units are logistic elements. Similar training curves were obtained for both monotonic and gaussian formats. Neither algorithm had serious problems with local minima (i.e., getting stuck at suboptimal solutions).[1]

Response properties

We studied the response properties developed by the hidden units during training in the same manner as Zipser and Andersen did for their model, except we plot the units' probability of firing (which is a continuous variable) instead of its instantaneous output (which is binary). The probability of firing can be thought of as equivalent to the firing rate, and thus equivalent to the continuous output in the Zipser and Andersen model. In other words, over a number of repeated presentations of a given input pattern, the frequency with which a binary unit's output is one encodes a continuous value, which can be conceived as a firing rate.

An interesting feature of area 7a neurons is that the visual and the eye position contributions to their overall response interact nonlinearly. For a constant retinal stimulus position, the total response is not composed of a constant visual response to which an independently varying amount of activity is added as the eye position changes. As Figure 7a and the work of Andersen and Zipser (1988; Zipser and Andersen, 1988) show, the visual and eye position components can vary simultaneously, in either the same or opposite directions, or to different degrees with eye position (see Andersen and Zipser, 1988, for a more detailed analysis of area 7a gain fields). When

1. The frequency of local minima was around 5% for back-propagation, and approximately 1% for the $A_{R\text{-}P}$ algorithm, in approximately 200 different simulations. One reason for the rather high frequency of local minima for backpropagation is likely the small number of hidden units in the network. The $A_{R\text{-}P}$ networks was less affected by this parameter, mainly because the unit's output noise improved the network's chances of escaping local minima.

examined after training, the hidden units of both types of $A_{R\text{-}P}$ networks displayed gain fields similar to those of area 7a neurons (Fig. 7b-c), as well as the same type of variety. The overall response of the hidden units, moreover (thin circles in Fig. 7), was always roughly planar along vertical and horizontal eye positions. This result was found in 78 percent of spatially tuned area 7a neurons (Andersen and Zipser, 1988). When a second hidden layer of binary stochastic units was been added to either the Mixed $A_{R\text{-}P}$ or All $A_{R\text{-}P}$ network, both networks learned to perform the task, and the units in the first hidden layer still developed planar gain fields similar to those of area 7a (Fig. 7d, only All $A_{R\text{-}P}$ shown).

It is worth noting that when studied in more detail, that is, when sampled over a larger range of eye positions, the gain fields produced by $A_{R\text{-}P}$ (as well as backpropagation) training are not exactly planar, but roughly sigmoidal along one direction of eye position (Fig. 8). In other words, the overall responses are approximately planar within a range of eye positions, and are flat outside this range. It turns out that this range is determined by the most eccentric eye positions on which the network was trained. The unit whose gain field is shown in Figure 8, for example, was part of a network trained with eye positions between -40 and 40 degrees (horizontally as well as vertically), and it developed a gain field approximately planar over this range along the x direction (there were other units in the network with similar gain fields oriented along the y direction). This result shows that the hidden units learn to interpolate for eye positions between those in the training set, but they do not learn to extrapolate to eye positions outside this range. We believe that this is a direct consequence of using a sigmoidal probability function (or input-output function in the case of the deterministic logistic element) for the hidden units. The gain fields of area 7a neurons may also flatten outside a certain range of eye positions, producing a sigmoidal shape like that in Figure 8. At present, the recording data available is too noisy, and perhaps too limited in range of eye positions, to distinguish between a simple plane and a sigmoid for the gain fields.

There was also a qualitative similarity between the visual receptive fields developed by the network's hidden units and those of area 7a neurons, as shown in Figure 9. The most striking feature of these neurons' receptive fields is their size, which extends to diameters of 80 degrees (Fig. 9a). This feature is reproduced by our model networks (Fig. 9b-c). Another feature is the complexity of these receptive fields' surfaces, characterized by one or more smooth peaks of various eccentricities, which sets area 7a neurons apart from those of many other visual areas. The networks' hidden units also display a similar complexity in their receptive fields, although such a comparison can be qualitative at best. As was the case for the gain fields, the addition of an extra hidden layer did not significantly affect the types of receptive fields developed by units in the first hidden layer (Fig. 9d, only All $A_{R\text{-}P}$ shown).

The response properties of the $A_{R\text{-}P}$ networks' hidden units are not only similar to those of area 7a neurons, but also to those of hidden units of networks trained by backpropagation to compute coordinate transformations. These response properties were described by Zipser and Andersen (1988; Andersen and Zipser, 1988). We were able to reproduce them also in a backpropagation network with the same number of units and the same training set as the $A_{R\text{-}P}$ networks (Fig. 10). This similarity suggests that the S'-model $A_{R\text{-}P}$ rule and backpropagation compute similar solutions (i.e., sets of network connection strengths) to the coordinate transformation problem.

Comparison of solutions

The solutions found by S'-model $A_{R\text{-}P}$ training and by backpropagation are not just similar in the qualitative sense depicted in Figures 7, 9 and 10. In fact, for a given training set, we found that the set of weights trained by the $A_{R\text{-}P}$ algorithm may be transferred to a backpropagation network (with continuous output hidden units) without any appreciable reduction in the accuracy of the network's response to that training pattern, and vice versa (Fig. 11). This is true for both versions of our networks (Mixed and All $A_{R\text{-}P}$) and for

networks with one and two hidden layers, as long as the output format is the same for the $A_{R\text{-}P}$ and backpropagation networks being compared. The individual values of the weights are *not* the same after training by the three different procedures, but the overall structure of these weights is such that the two algorithms' solutions to the coordinate transformation problem are functionally equivalent for the various network structures.

Generalization

A property that is often exhibited by artificial neural networks trained by examples is the ability to generalize from those examples, that is, to produce the correct output when presented with input patterns that were not in the training set. This property is of great theoretical and practical importance, as it demonstrates that the model network has not merely learned to associate patterns in the training set with their correct outputs on an individual basis, but has learned to perform the transformation implied in the training examples. In our task, this mapping is the addition of two position vectors.

We tested two extensively trained networks for two types of generalization abilities. One is the ability to perform the correct vector addition of new, random input patterns that code for the same output locations as the training set. As shown in Figure 12 (*left*), all three networks performed this task extremely well. The other generalization task required the trained networks to give the correct output for input patterns coding for new output locations (Fig. 12, *right*), which is a more difficult task. Although all networks produced some error, this was still considerably less than for the untrained nets, indicating that the networks generalized to a considerable extent.

DISCUSSION

Choice of the learning algorithm

The choice of the algorithm used to update the network's connection strengths was the central issue of our study. There exist a number of

procedures to change the weights of a network in order to maximize some measure of performance in a supervised learning paradigm (for review, see Anderson and Rosenfeld, 1988; Barto, 1985; Hinton, 1987; Lippmann, 1987; McClelland and Rumelhart, 1988). An important class of such learning algorithms consists of those that maximize the performance measure by following its gradient (i.e. the direction of its maximum increase) with respect to the weights and adjusting them accordingly, arriving at the set of weights that produces the optimal output for every pattern in the training set. Backpropagation is an important and powerful algorithm belonging to this class. It has been used to train networks to perform such disparate tasks as pronouncing written English text (Sejnowski and Rosenberg, 1986), to detecting explosives in passengers' luggage at airports (Shea, 1989). It is also the training procedure used by Zipser and Andersen to teach their network to perform the coordinate transformation task (Zipser and Andersen, 1988). Powerful as it is, however, backpropagation suffers the problem mentioned above of not being easily implementable in biological neuronal hardware. Central to the backpropagation algorithm are: i) the feedback of detailed error signals that are specific for each output unit; ii) the retrograde propagation of these signals from the output units back to the hidden and input units; and, iii) the adjustment of synaptic strengths using global network information, that is, information about the activities and errors of units removed from the synapse whose strength is adjusted. These requirements represent major hurdles to envisioning backpropagation as a plausible model of learning in biological neural networks even to only a rough approximation, a concern that has been expressed by some of the discoverers of this algorithm (Rumelhart et al., 1986b). Possible solutions to this problem have been proposed, which require, for example, specialized connections carrying the error signals for each unit in the network (Hecht-Nielsen, 1989), or symmetrical feedback pathways with connection strengths identical to those in the forward network (Parker, 1985; Zipser and Rumelhart, 1990). Besides being rather unconvincing from a biological perspective, however, these

complicated mechanisms detract quite a bit from the simplicity of structure and process that makes parallel neural networks so appealing as models of nervous system function.

We chose the $A_{R\text{-}P}$ learning algorithm because it uses the same abstract principles involved in supervised learning as backpropagation, but with specific processes that are more plausible for implementation in neurobiological hardware. In particular, in the $A_{R\text{-}P}$ algorithm: i) the feedback signal is a single scalar value computed from the output units' average error; ii) this signal is the same for all units in the network, and as such it does not require backwards propagation along network connections; and iii), synaptic strength is adjusted using the payoff signal and information about the activity of the presynaptic and postsynaptic unit only. We will return to these three features shortly.

How the network learns

The abstract principle used by backpropagation is *gradient descent*—the minimization of an error measure by following the negative of its gradient with respect to the weights. A priori there are no conceptual or experimental obstacles to envisioning the brain using this principle too, given a plausible performance measure. The S'-model $A_{R\text{-}P}$ algorithm, as implemented in both of our network classes, also makes use of this general principle. While the backpropagation algorithm, however, computes the exact value of the error's gradient for a given input pattern, the S'-model $A_{R\text{-}P}$ rule computes only an estimate of that gradient (Barto and Jordan, 1987; Williams, 1986, 1987). Units trained by the $A_{R\text{-}P}$ rule do not have the detailed information about the error vector and the state of other units which is necessary to compute the exact gradient and which backpropagation units obtain through non-biological pathways. Due to the random noise in their output, however, $A_{R\text{-}P}$ units can "jitter" their activity during learning so as to get an estimate of how the noise in activity affects the payoff they receive, which in turn allows them to estimate the direction in weight space along which to change their weights in order to increase reinforcement.

While this method allows $A_{R\text{-}P}$-trained units to properly adjust their weights using only locally available information, it is more random in its search for a solution than backpropagation. These differences are obvious in the learning curves of Figure 6. Backpropagation's precise computation of the performance gradient tells the algorithm the exact manner in which to change the weights so that the error is monotonically decreased, resulting in the smooth curve of Figure 6a. As this curve shows, the error falls quickly to a value below the resolution of the gaussian units in the retinal input (10 degrees), and then continues to decrease much more slowly as the output is refined to match the training signal. The curves for $A_{R\text{-}P}$ learning (Fig. 6b-c) follow an envelope very similar to the backpropagation curve, but they are much noisier. The noise is due to the randomness of the $A_{R\text{-}P}$ units' output. In order to learn, the $A_{R\text{-}P}$ element adjusts its weights so that its net input drives its probability of firing towards one of the flat regions of the sigmoid function, thus effectively decreasing the randomness of its output. The decrease of the $A_{R\text{-}P}$ units' jitter as learning proceeds is reflected in decreasing noise on the learning curves.

Biological plausibility of the $A_{R\text{-}P}$ algorithm

As we mentioned above, one crucial requirement for our choice of a learning algorithm was a greater plausibility of biological implementation than the backpropagation algorithm. We must point out at the outset, however, that $A_{R\text{-}P}$ networks were not designed as literal models of biological neural nets (Barto, 1985, 1989). Because of the poor knowledge we have of the precise mechanisms of information processing used by nervous systems, the most useful connection between artificial and biological neural networks is presently limited to the description of abstract processes in simplified models and the investigation of the *possibility* of implementation in the biological hardware. In other words, the $A_{R\text{-}P}$ element was not designed by collecting scattered known facts of neurobiology and molding them into a computationally interesting unit capable of supervised learning, but

rather as a simple, "neurally inspired" element with a few theoretically motivated features that give it interesting learning abilities. We will discuss biological plausibility, therefore, in its literal sense of suggesting that the abstract computing processes performed by the $A_{R\text{-}P}$ unit during learning are more in keeping with possible neural mechanisms proposed and partly demonstrated by experimental neuroscientists than the mechanisms used by backpropagation networks.

The first and most salient element of $A_{R\text{-}P}$ models which aligns them with many neurobiological and psychological models of learning is the scalar payoff signal, r. This has the attractive features of being computed from an average value of the error of the output units, and of being transmitted as a single value to all the $A_{R\text{-}P}$-trained units in the network equally. This error could also be detected, for example, as a function of the angular difference between an object in space and the end position of a reaching arm movement or a saccade toward that object. After successful training, this difference would be nil and reinforcement would be maximal. The reinforcement signal could thus be computed by a part of the nervous system that monitored the animal's behavior with very little information about the activity of area 7a neurons. The fact that a single value is valid for All $A_{R\text{-}P}$ units implies that only one connection is necessary from the reinforcement computing region to area 7a. In the backpropagation algorithm, on the other hand, the output units' errors must be kept as separate components as they are fed back to the network to adjust individual weights.

A single scalar value is easier to transmit to a group of neurons than an error signal with specific multiple components. Evidence that the nervous system may use such signals already exists. For example, the nucleus basalis sends a widespread system of cholinergic connections to several cortical areas, and the signals involved appear to be related to behavioral choices and reward (Richardson et al., 1988). The payoff signal required by our model, of course, would not have to be distributed on such a wide scale. The signal could carry

information about a specific motor task, for example, the accuracy of a saccade to a target, and thus be used only by one or a few ensembles of neurons in area 7a. Because a single signal, however, would be valid for an entire group of neurons, there would be no need to propagate it through special pathways to specific units in the network of interest. This feature of the $A_{R\text{-}P}$ algorithm is more attractive than backpropagation's requirement for the retrograde propagation of error signals along specific pathways.

Besides not having to carry specific information to train individual neurons, the payoff signal used in our networks has the advantage that it can be independent of any coordinate system. In backpropagation, the "teacher" signal must code for the correct head-centered location as a vector in a craniotopic reference frame. The $A_{R\text{-}P}$ algorithm, on the other hand, computes its feedback signal from the average of the output error's absolute value (Eqs. 2-3), which is a single number that can be derived from the comparison of retinotopic as well as craniotopic positions.

Another "biological" feature of learning by $A_{R\text{-}P}$ units is the use of information that is locally available to the element itself at its individual synapses, in a fashion reminiscent of Hebbian learning. The weight-adjusting equation for the i^{th} $A_{R\text{-}P}$ unit (Eq. 1) consists of the sum of two terms, each assigning the "reward" part and the "penalty" components, respectively, of the learning rule. These terms consist of three components:

i) the payoff signal, r (and the corresponding penalty value, $1\text{-}r$);
ii) information regarding the current state of the unit $(x_i\text{-}p_i)$;
iii) the output of the presynaptic element, x_j, directly available at the synapse that the j^{th} unit makes onto the i^{th} unit.

We have already discussed (i). In (ii), x_i is the unit's output (zero or one), and p_i is the probability that the unit's output will be one given the current net input, which depends on the unit's weights. As mentioned above, this probability could also be interpreted as the rate at which the unit will fire given the present input. These two values, as well as x_j (iii), are effectively available at the connection between the input unit and the given hidden unit. The $A_{R\text{-}P}$ rule therefore embodies one of the most important elements of Hebbian learning (Hebb, 1949), that is, the proportionality of a change in synaptic strength to both presynaptic and postsynaptic signals. Hebbian learning remains one of the most attractive mechanisms for synapse modification, both on theoretical (Linsker, 1989) and experimental grounds (Ito, 1984; Kelso et al., 1986; Sejnowski et al., 1989; Stanton and Sejnowski, 1989; Brown et al., 1990). This is in contrast to backpropagation, in which changes in strength at one connection require information about the activities and error signals at all the connections in every layer above that connection (Rumelhart et al., 1986a).

The Mixed $A_{R\text{-}P}$ network, as we have mentioned, uses the $A_{R\text{-}P}$ rule only to train its hidden units. Its output units are trained by the delta rule (Eq. 4). Although this is the same rule as is used for the output units in backpropagation training, this does not pose as many obstacles to biological implementation as the full backpropagation algorithm does. As shown in Figure 5a, the only extra information required by the delta rule, as compared to the $A_{R\text{-}P}$ rule, is an error vector from the external evaluator. This is needed to form an individual error signal for each output unit, whose weights are then adjusted by error correction. There is no requirement, however, for backpropagation of error or activity signals across synapses. In fact, the delta rule also has a Hebbian form at the output layer, again in the sense that all the information required to adjust a connection's strength is available at the synapse. In Eq. 4 two terms are multiplied, the first of which (in square brackets) contains "postsynaptic" information, while the other is the activity of the "presynaptic" unit.

The last feature that adds some biological flavor to the $A_{R\text{-}P}$ unit is the probabilistic nature of its output. The unpredictability of the exact firing rate produced by a neuron for any given presentation of a certain input has long been recognized as a feature of nerve cells (see, for example, Vogels, Spileers and Orban, 1989; Tolhurst, 1989; Tolhurst et al., 1983). In fact, this stochastic aspect of activity is one of the reasons neurophysiologists usually present data as summed histograms of several trials (Sejnowski, 1981). This is a feature that is not included in the deterministic units of backpropagation networks. The binary stochastic element's output is not simply noisy. It has a variance that is an increasing function of the mean probability of firing. The variability of spike trains recorded from cortical neurons also exhibits this statistical property (Vogels et al., 1989).

Many discussions of this aspect of neural activity have emphasized the difficulties it creates, such as the requirement it may impose on certain types of sensory information to be distributed over populations of neurons (Tolhurst, 1989). In our model, however, the intrinsic variability of the $A_{R\text{-}P}$ units' responses to input signals is essential for the learning process. It allows the network to jitter its weights around the current set of values, thus sustaining the search for a better solution. The noise provides the algorithm with the means of obtaining information about local variations in reinforcement. By making successive incremental adjustments to the weights, the algorithm converts these local variations into an estimated gradient of the reinforcement signal. In this manner the noise compensates, in a sense, for the scarcity of information contained in the scalar payoff signal. The stochastic aspects of the model, therefore, are not mere demonstrations of robustness to noise. The $A_{R\text{-}P}$ rule demonstrates, rather, how a computational unit's output variance can be used to achieve learning in a network that receives less than optimal feedback information.

Simulation results

The basic results of this study corroborate the validity, from a physiological perspective, of parallel networks with distributed representations as models of area 7a. We have shown that the $A_{R\text{-}P}$ algorithm can train a neural network to perform the same coordinate transformation task as that performed by Zipser and Andersen's model. We also found that the solution discovered by this algorithm is equivalent to that found by backpropagation. As was established in Zipser and Andersen's analysis (Zipser and Andersen, 1988), this solutions give hidden unit response properties (planar gain fields and large visual receptive fields) very similar to those of area 7a neurons presumed to code for spatial location. These response properties, therefore, are not a specific result of the backpropagation training procedure. The set of connections strengths computed by the $A_{R\text{-}P}$ algorithm, moreover, is not a unique one imposed by the network's architecture. Other solutions, not involving planar gain fields or large receptive fields, can be constructed which would work for the training sets we used. It is striking, then, that $A_{R\text{-}P}$ and backpropagation produce this particular algorithm for computing coordinate transformations, and not any other.

In a more detailed analysis of the model we have shown that a second layer of hidden units can be added to the network without changing the response properties of the first hidden layer, and that the model networks are indeed capable of generalizing their coordinate transformation abilities to new input patterns. Both these results strengthen the physiological significance of this model architecture. The former implies that relay elements—an important and ubiquitous feature of brain architecture—are not an obstacle to learning, and allow similar solutions to develop. The latter establishes the important point that these model networks are indeed learning to perform the coordinate transformation task. They do not merely act as content-addressable memories—associating each input pattern in the training set with its correct output individually— but rather they are capable of

abstracting from the training examples the transformation common to them, in this case vector addition, and applying it successfully to new pairs of retinal and eye positions. This property has been observed before in parallel networks with very few hidden units in the hidden layer compared to the input layer (Cottrell et al., 1987).

The number of training iterations required for convergence by $A_{R\text{-}P}$ and backpropagation were comparable for networks and training sets of the size we used. We have not examined in our study the issue of how the $A_{R\text{-}P}$ algorithm behaves for networks with considerably larger numbers of hidden units and training locations. From previous experience with this learning rule (Barto and Jordan, 1987), learning should be significantly slower for such networks. It may be possible, however, to make the algorithm more resistant to scale changes, for example, by using a topographically more specific reinforcement signal. Our use of a single scalar feedback signal could thus be viewed as a worst-case scenario that does not exclude more specialized signals which may be used by biological systems.

Future directions

It would be desirable to modify the $A_{R\text{-}P}$ algorithm so that it could train networks with continuous-output hidden units. Actually, any algorithm capable of performing gradient descent using a scalar reinforcement signal would be acceptable. Gradient-descent algorithms that use scalar payoff signals are currently being developed (e.g. Gullapalli, 1988), and it would be a natural continuation of this work to try to apply them to networks modeling area 7a. The major hurdles in these algorithms involve the theoretical details of simultaneously updating the mean and variance of multi-parametric distributions required by continuous stochastic units. The present form of these algorithms is similar to that of the $A_{R\text{-}P}$ procedure for binary units. It is conceivable that the extension of the concepts of supervised $A_{R\text{-}P}$ learning to networks with continuous output units will be a natural refinement which should not drastically change the types of solutions obtained.

Conclusions

We have shown that: i) the $A_{R\text{-}P}$ algorithm can train neural networks to compute coordinate transformations; ii) the convergence times for small networks are comparable to those obtained by backpropagation training; iii) in the process of learning this computation, the hidden units develop gain fields and receptive fields qualitatively similar to those of area 7a neurons; iv) the solutions are equivalent to those computed by backpropagation; and, v) these networks generalize appropriately. We have also pointed out a number of features of the $A_{R\text{-}P}$ algorithm that bring it closer than backpropagation to what is known about biological learning. We must emphasize again, however, that the focus of our interest at this point is not in how literally $A_{R\text{-}P}$ networks reproduce individual neurophysiological processes. It is rather the fact that these algorithms form a family of training procedures that yield similar functional representations when applied to a class of parallel distributed networks, and that they can do so using mechanisms not excluded, and perhaps suggested, by neurophysiological evidence.

These results represent a step toward establishing the physiological validity of the architecture and general learning principles of the model of area 7a introduced by Zipser and Andersen. They show that physiological properties can arise from a more plausible learning algorithm than backpropagation, thus suggesting that the detailed processes by which neuronal ensembles learn may play only a secondary role in their ultimate collective behavior. Abstract optimization principles, such as gradient descent, may instead be more important determinants of neuronal learning strategies, and it would be worthwhile to pursue such hypotheses with further theoretical and experimental studies.

FIGURE LEGENDS

Figure 1:

a) Experimental method of measuring spatial gain fields of area 7a neurons. These experiments were carried out several years before our modelling project (Andersen, Essick and Siegel, 1985, 1987). The monkey faces a projection screen in total darkness, and is trained to fixate on a point, *f*, at one of nine symmetrically placed locations on a projection screen with his head fixed. The stimulus, *s*, is always presented at the same retinal location, at the peak of the retinal receptive field, *rf*. The stimulus consists of 1- or 6-degree diameter bright spots flashed for 500 milliseconds. *b)* Peri-stimulus time histograms of neuronal activity recorded from a particular area 7a neuron, arranged in the same relative positions as the corresponding fixation spots. The arrows indicate the time of visual stimulus onset. The characteristics of the response to the visual stimulus at the various angles of gaze constitutes the neuron's eye position gain field. (*a* and *b* adapted from Andersen, Essick and Siegel, 1985). *c)* Graphic representation of the gain field in *(b)*, introduced by Zipser and Andersen (1988). The diameter of the thin outer circle is proportional to the total response evoked by the stimulus. The annulus' diameter represents the contribution to the total response due to eye position alone, and is measured as the background activity recorded 500 ms before the stimulus onset. The dark inner circle represents the visual contribution to the response, and its diameter is computed by subtracting the background activity from the total response. This representation shows that this neuron's gain field is roughly planar in a direction up and to the left.

Figure 2:

a) Structure of the Mixed $A_{R\text{-}P}$ network. The network is composed of three layers of computing units: input units (encoding retinal location of stimulus and eye position), hidden units, and output units (encoding

head-centered location of visual stimulus). Retinal position of the visual stimulus is encoded topographically by an 8 x 8 array of input units, each with a gaussian receptive field (described in *(b)*). The remaining 32 input units code for eye position in a linear fashion (see *(c)*), with 2 groups of 8 units encoding horizontal gaze angle (with positive and negative slopes), and 2 groups of 8 units for vertical angle. Units in the output layer code for head-centered coordinates in a monotonic format *(output 1)* similar to the eye position input, or in a gaussian format *(output 2)* similar to the retinal input. The hidden units are binary stochastic elements, while the output units are deterministic logistic elements (see Fig. 3). *b)* Angle-coding function of the retinal input units. Each unit's activity level is a gaussian function of the retinotopic x and y coordinates of the visual stimulus, with $1/e$ width of 15 degrees, and spaced 10 degrees apart from that of its neighboring units. *c)* Angle-coding function for the eye position units. Each unit codes for the x or y eye position linearly. The slopes and intercepts for each unit were assigned randomly within ranges observed for area 7a neurons.

Figure 3:

a) Binary stochastic element. This neurally-inspired computing element performs a weighted sum of its synaptic inputs (x_0 through x_m) by multiplying each input by a synaptic weight (w_{i0} through w_{im}), which can be positive or negative, and adding these products. This sum (s_i) is then used to compute a value (p_i) between zero and one from the logistic function ($1 / (1 + exp\ (-s_i\ ))$). The element then produces an output of 1 with probability p_i, and an output of zero with probability $1\text{-}p_i$. We used this element in the hidden layers of all our networks, as well as in the output layer of the All $A_{R\text{-}P}$ network. *b)* Deterministic logistic element. This unit computes a weighted sum of its input in the same manner as the binary stochastic element. This sum is also passed through the logistic function, but the resulting value (x_i) is the unit's output itself. The unit therefore produces a

continuous output between zero and one which is determined exactly by the weighted sum of its inputs. This is the element that was used in the hidden and output layers of the Zipser and Andersen model. We used it only in the output layer of the Mixed $A_{R\text{-}P}$ network.

Figure 4:

a) Structure of the All $A_{R\text{-}P}$ network. The input and hidden layers are the same as for the Mixed $A_{R\text{-}P}$ network (Fig. 2). The output layer is composed of binary stochastic units (Fig. 3a). These code for locations in craniotopic space by dividing the latter into discrete regions according to one of two formats (output 1 and output 2) described in *(b)* and *(c)*. *b)* Binary-monotonic format for the All $A_{R\text{-}P}$ network. Each unit transforms an output value between zero and one into and angle via a step function. An output of zero corresponds to all angles less (positive step) or greater (negative step) than a given "cut-off" angle, and an output of one codes for all angles greater or less, respectively, than the cut-off angle. We use step here to indicate the direction along which the step function changes from zero to one. Typically there are four sets of three units each, for horizontal and vertical coding with positive and negative step. Within each set, the cut-off angles for the three units are at -40, 0 and 40 degrees. Only the functions for one positive-step triplet of units are plotted. *c)* Binary-gaussian format. Four units divide craniotopic space into 13 regions using overlapping circular binary receptive fields. Each unit outputs a 1 for a position within a 100-degree radius circle centered at one of the four positions ($\pm 60, \pm 60$) degrees.

Figure 5:

a) Learning scheme for the Mixed $A_{R\text{-}P}$ network. Training proceeds in two phases that are repeated sequentially. First, a pair of retinal and eye positions is presented at the input layer. The signal propagates forward (*solid arrows*) to the network's upper layers along connections strengths that initially have random or zero values. The network's

output is evaluated by some external agent, and two types of signals are fed back to the network (*dashed arrows*). One is a vector error signal that consists of the differences between the actual and desired outputs for the output units, and is sent to the output units. The other is a scalar payoff signal (r) between zero and one that is sent to the hidden units. In the second phase the connection weights are adjusted using the error and payoff signals. The output units adjust their weights according to the delta rule, while the hidden units adjust them by the S'-model $A_{R\text{-}P}$ learning rule. The backpropagation network used for reference was trained by the standard backpropagation algorithm described by Rumelhart et al. (1986a). *b)* Learning scheme for the All $A_{R\text{-}P}$ network. This is identical to Mixed $A_{R\text{-}P}$ learning described in *(a)*, except that the scalar payoff signal r is broadcast to all the units in the hidden and output layers, and all the weights are adjusted by the S'-model $A_{R\text{-}P}$ rule. The network therefore employs only the scalar payoff signal for feedback information on its performance, and no error vector is required.

Figure 6:

Learning curves for the various networks studied. The error plotted is the absolute value of the difference between the network's expected and actual output, averaged over the units in the output layer and over the patterns in the training set. This average error corresponds approximately to the radial distance between the craniotopic location encoded by the output layer and the correct one (given by the sum of the retinal and eye position vectors). The dashed line is a scaling reference, 10 degrees corresponding roughly to the resolution of the visual input. A three-layer network architecture with 3 hidden units was used in *a-c*. The training set consisted of 12 random inputs coding for 4 spatial locations. A two-unit monotonic output format was used, which provided for easy conversion of the error values from unit activation levels to angular coordinates of craniotopic space. The training set was chosen small for better visualization of network

behavior. The error for the binary output units of the All $A_{R\text{-}P}$ network was converted to degrees using the same linear activation function as for the other two networks. *a)* Backpropagation training. *b)* Mixed $A_{R\text{-}P}$ training. *c)* All $A_{R\text{-}P}$ training.

Figure 7:

Eye position gain fields for area 7a neurons and for the model networks' hidden units. Gain fields for four different neurons *(a)* or hidden units *(b-d)* are shown in each case. The nine circles in each box are a set of responses sampled at nine different eye position, with the stimulus' retinal location held constant. As described in Figure 1, the thin outer circles represent the total activity (normalized), the dark inner circles are proportional to the visual stimulus' contribution to the total response, and the white annuli are the background activity due to eye position. By "activity" we mean frequency of firing for area 7a neurons, and probability of firing for the networks' hidden units. The spacing between eye positions is twenty degrees for area 7a neurons. It is twenty degrees for all the networks' hidden units, except for the two gain fields in the bottom left of *(c)* and *(d)*, for which the spacing is forty degrees. *a)* Area 7a neurons. *b)* Mixed $A_{R\text{-}P}$ network. *c)* All $A_{R\text{-}P}$ network. *d)* All $A_{R\text{-}P}$ network with two layers of hidden units.

Figure 8:

Probability of firing (total response alone) of a hidden unit from a Mixed $A_{R\text{-}P}$ network, sampled over a continuous range of eye positions. The gain field is planar over a wide range of eye positions. Note the slight saturation effects (flattening) at the edges of the eye position field.

Figure 9:

Visual receptive fields of area 7a neurons and of networks' hidden units. As in Figure 7, the variables sampled are firing rate for area 7a neurons *(a)* and probability of firing for the hidden units *(b-d). a)*

Area 7a neurons' receptive fields. Each was sampled at seventeen points in a forty-degree-radius circle, and a smooth surface was obtained by gaussian interpolation (adapted from Zipser and Andersen, 1988). *b)* Mixed $A_{R\text{-}P}$ network. *c)* All $A_{R\text{-}P}$ network. *d)* All $A_{R\text{-}P}$ network with two layers of hidden units.

Figure 10:

Response properties of hidden units in a backpropagation-trained network. Shown are representative gain fields *(a)* and the receptive fields *(b)* recorded from the hidden units of a backpropagation after it was trained to perform the coordinate transformation. Note the similarity between these response properties and those of the hidden units of $A_{R\text{-}P}$-trained networks (Figs. 7 and 9). Two of the receptive fields in *(b)* (top right and middle one on the left) are adapted from Zipser and Andersen (1988).

Figure 11:

Error values (as defined for Fig. 6) produced by various networks when trained sets of connection strengths were swapped among them. Three networks (backpropagation, Mixed Ar-p and All Ar-p) were first trained to a given accuracy. The trained weights were then exchanged among the different networks, and the error upon presentation of the training set was recorded. The value of ten degrees is again used as a reference (see Fig. 3). The average error for all untrained networks was around 60 degrees. The error values do not show a significant change when backpropagation training is replaced by Mixed or All Ar-p training, and vice versa, showing that the solutions found by the different algorithms are functionally equivalent. BP = backpropagation. *a)* Backpropagation trained weights tested on backpropagation network. *b)* Mixed $A_{R\text{-}P}$ trained weights tested on backpropagation network. *c)* Mixed $A_{R\text{-}P}$ trained weights tested on Mixed $A_{R\text{-}P}$ network. *d)* Backpropagation trained weights tested on Mixed $A_{R\text{-}P}$ network. *e)* Backpropagation trained weights tested on

backpropagation network (binary output format of All $A_{R\text{-}P}$ network used in the training set). *f)* All $A_{R\text{-}P}$ trained weights tested on backpropagation network. *g)* All $A_{R\text{-}P}$ trained weights tested on All $A_{R\text{-}P}$ network. *h)* Backpropagation trained weights tested on All $A_{R\text{-}P}$ network.

Figure 12:

Generalization properties. Three-layer networks were trained by the three different algorithms on a set of forty random input pairs. The nets were then tested on a new pattern set, and the average absolute error (as defined for Fig. 6) was recorded (*solid bars*). The error for this testing set was also recorded for the untrained networks for reference (*hatched bars*). In the test for the recognition of the training output locations *(left)*, the testing set consisted of forty new random inputs that coded for the same four spatial locations as in the training set. This tested whether the nets had really learned to add the eye position and retinal location vectors to obtain the resulting craniotopic location, and not just formed an associative memory storage of the training set. In the test for generalization to new locations *(right)*, the testing set consisted of random inputs that coded for 40 new random head-centered locations. This tested for the networks' ability to generalize the vector addition operation to new targets. Note that the error bars labelled Mixed $A_{R\text{-}P}$ were obtained (for the second task only) by transferring a set of Mixed $A_{R\text{-}P}$ trained weights on a deterministic network and testing for generalization using this network. The reason for doing this was that the Mixed $A_{R\text{-}P}$ network cannot perform the second generalization task because there are too few hidden units to produce new locations in the continuous output format. The All $A_{R\text{-}P}$ network does not have this problem because the binary output format codes for regions of space and not for unique locations. BP = backpropagation.

a
FIX CENTER
FIX LEFT
b
ALL STIM. RETINAL (20.-20)
FIX: -20,20
0,20
20,20
-20,0
0,0
20,0
-20,-20
0,-20
20,-20
c

Figure 1

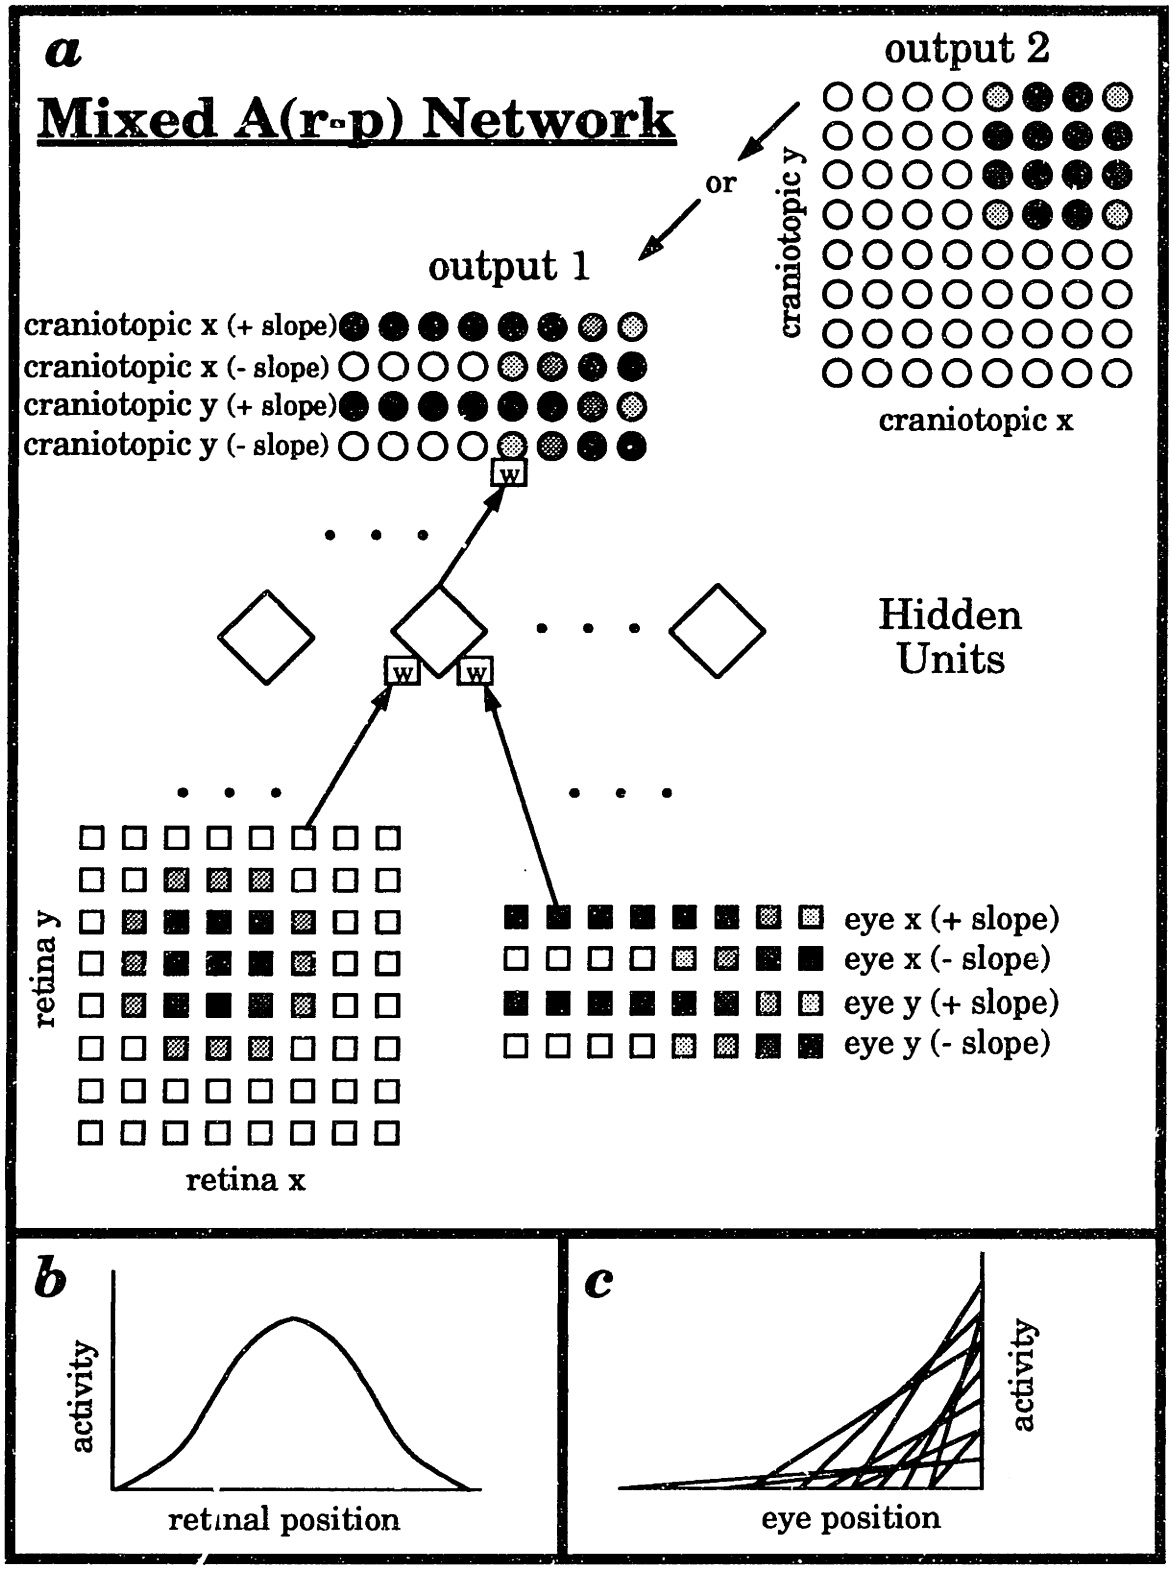
a
Mixed A(r-p) Network
output 2
or
craniotopic y
output 1
craniotopic x (+ slope)
craniotopic x (- slope)
craniotopic y (+ slope)
craniotopic y (- slope)
craniotopic x
w
Hidden Units
retina y
retina x
eye x (+ slope)
eye x (- slope)
eye y (+ slope)
eye y (- slope)
b
activity
retinal position
c
activity
eye position

Figure 2

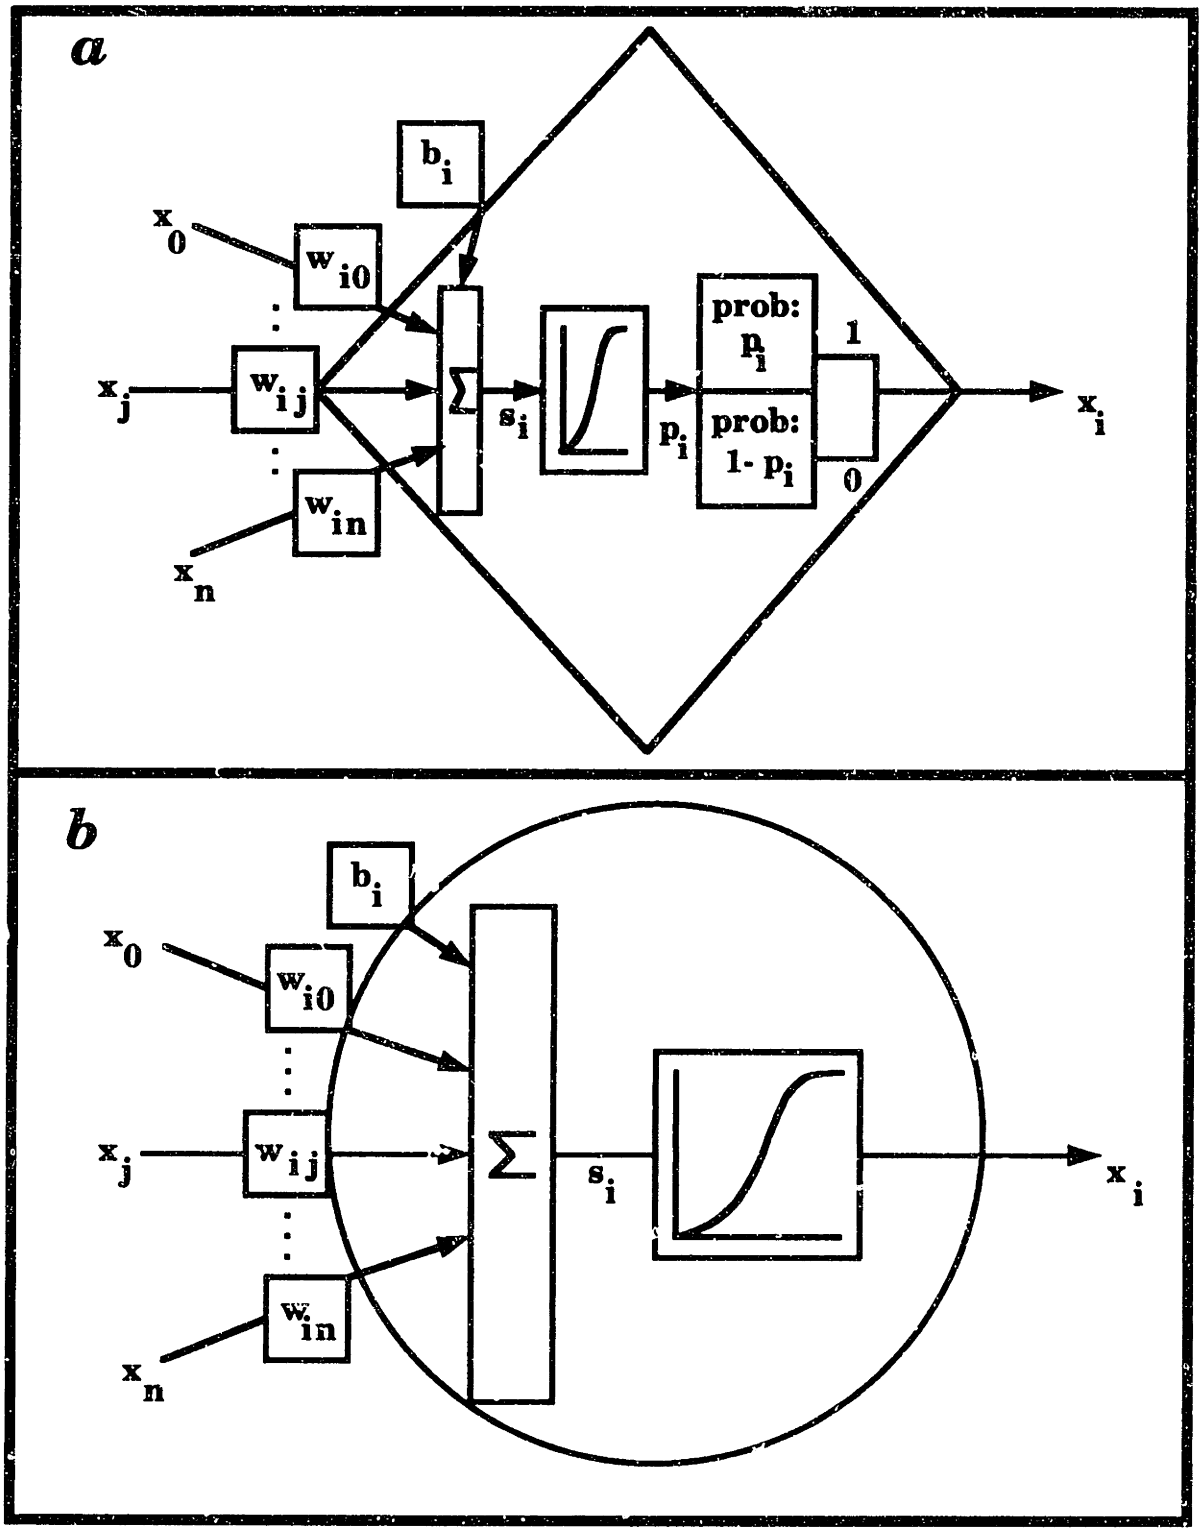
a
b_i
x_0
w_i0
w_ij
x_j
w_in
x_n
Σ
s_i
p_i
prob: p_i
prob: 1- p_i
1
0
x_i
b
b_i
x_0
w_i0
w_ij
x_j
w_in
x_n
Σ
s_i
x_i

Figure 3

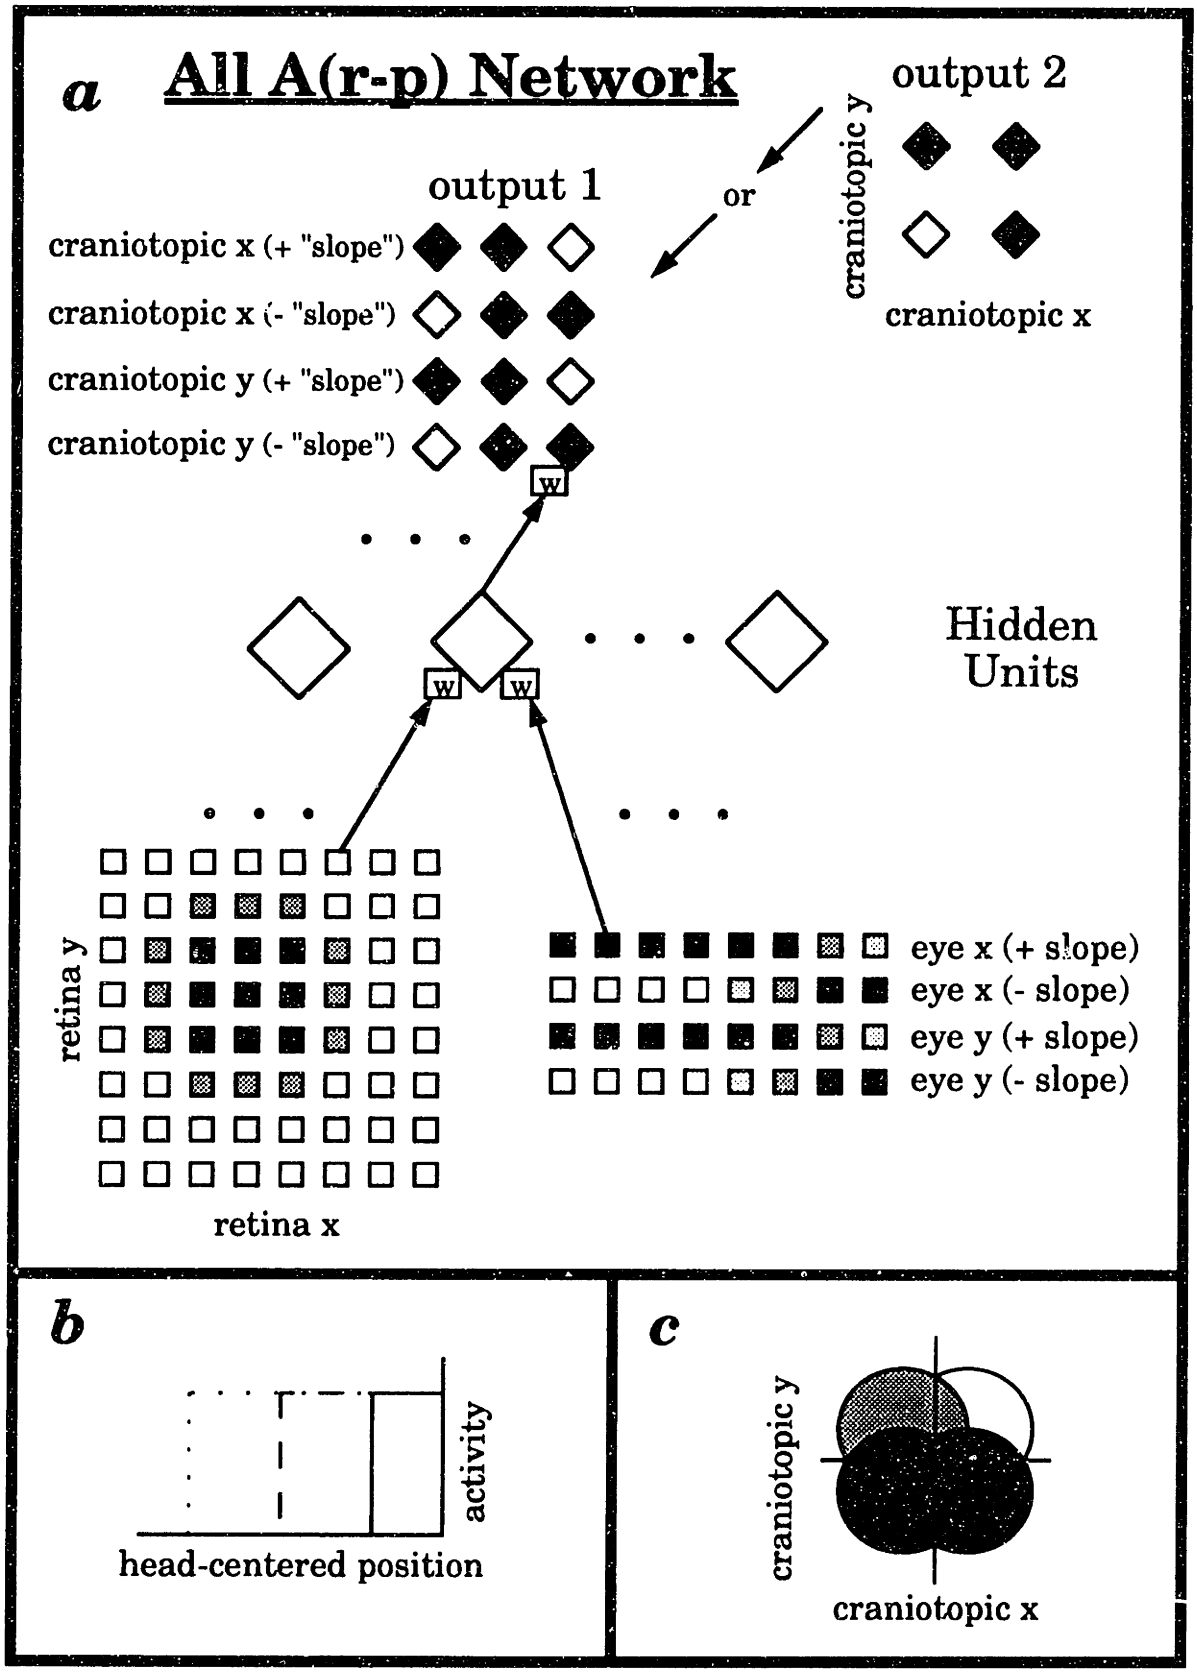
a
All A(r-p) Network
output 2
craniotopic y
craniotopic x
or
output 1
craniotopic x (+ "slope")
craniotopic x (- "slope")
craniotopic y (+ "slope")
craniotopic y (- "slope")
w
Hidden Units
retina y
retina x
eye x (+ slope)
eye x (- slope)
eye y (+ slope)
eye y (- slope)
b
activity
head-centered position
c
craniotopic y
craniotopic x

Figure 4

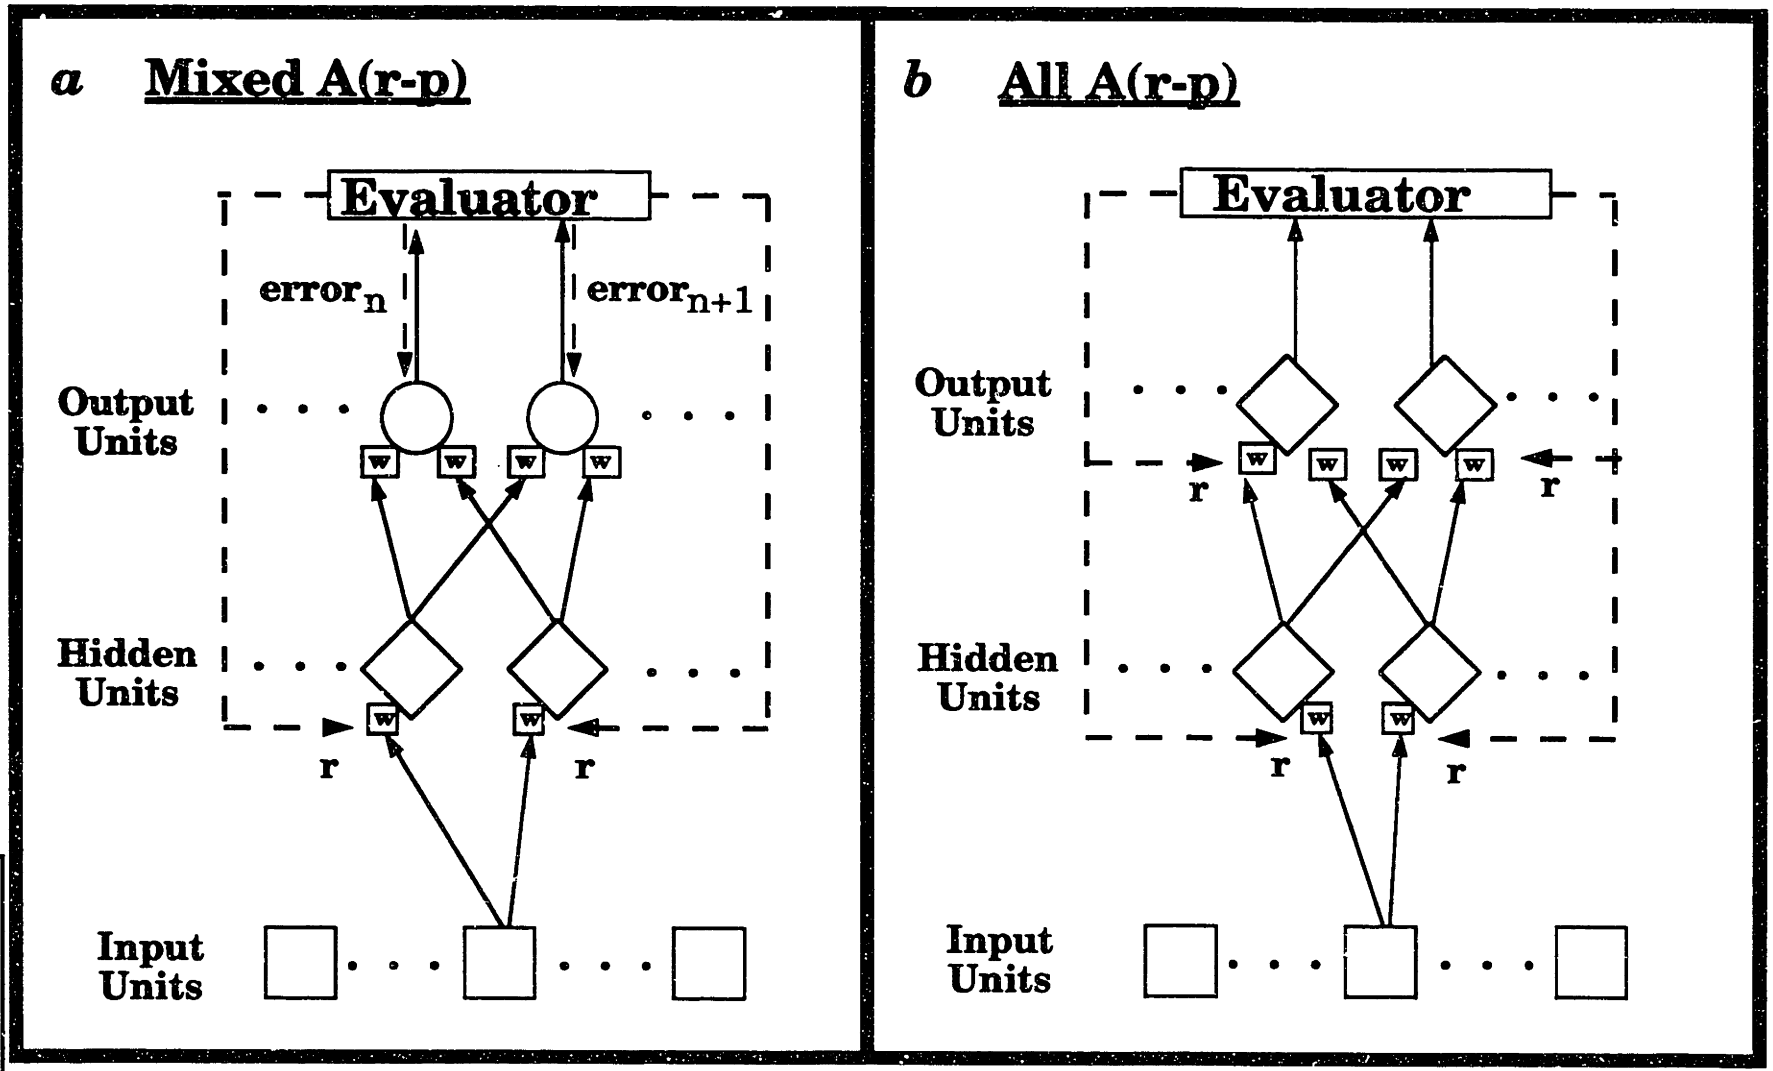
a Mixed A(r-p)
Evaluator
error$_n$
error$_{n+1}$
Output Units
Hidden Units
Input Units
w
r
b All A(r-p)
Evaluator
Output Units
Hidden Units
Input Units
w
r

Figure 5

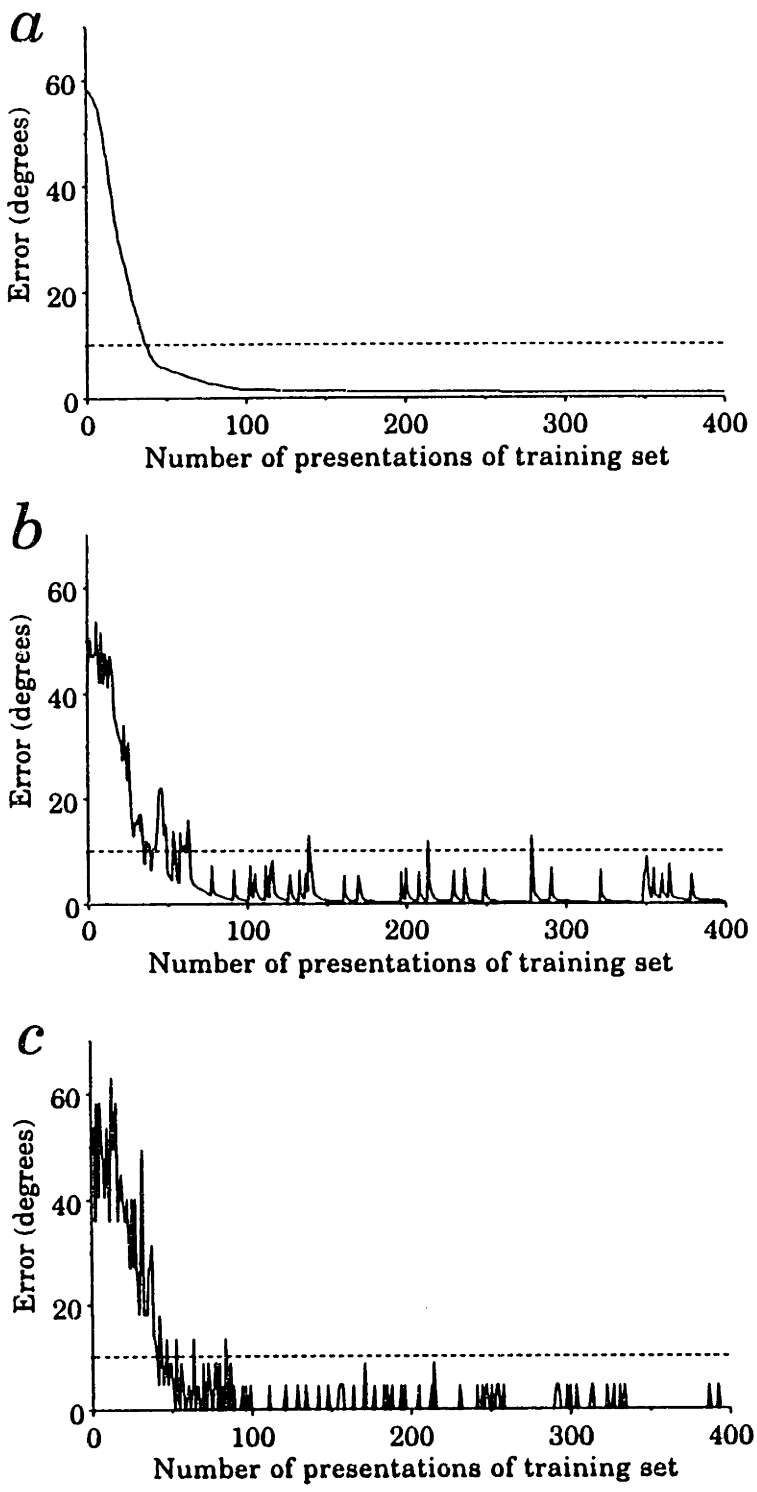
a
60
40
20
0
Error (degrees)
0
100
200
300
400
Number of presentations of training set
b
60
40
20
0
Error (degrees)
0
100
200
300
400
Number of presentations of training set
c
60
40
20
0
Error (degrees)
0
100
200
300
400
Number of presentations of training set

Figure 6

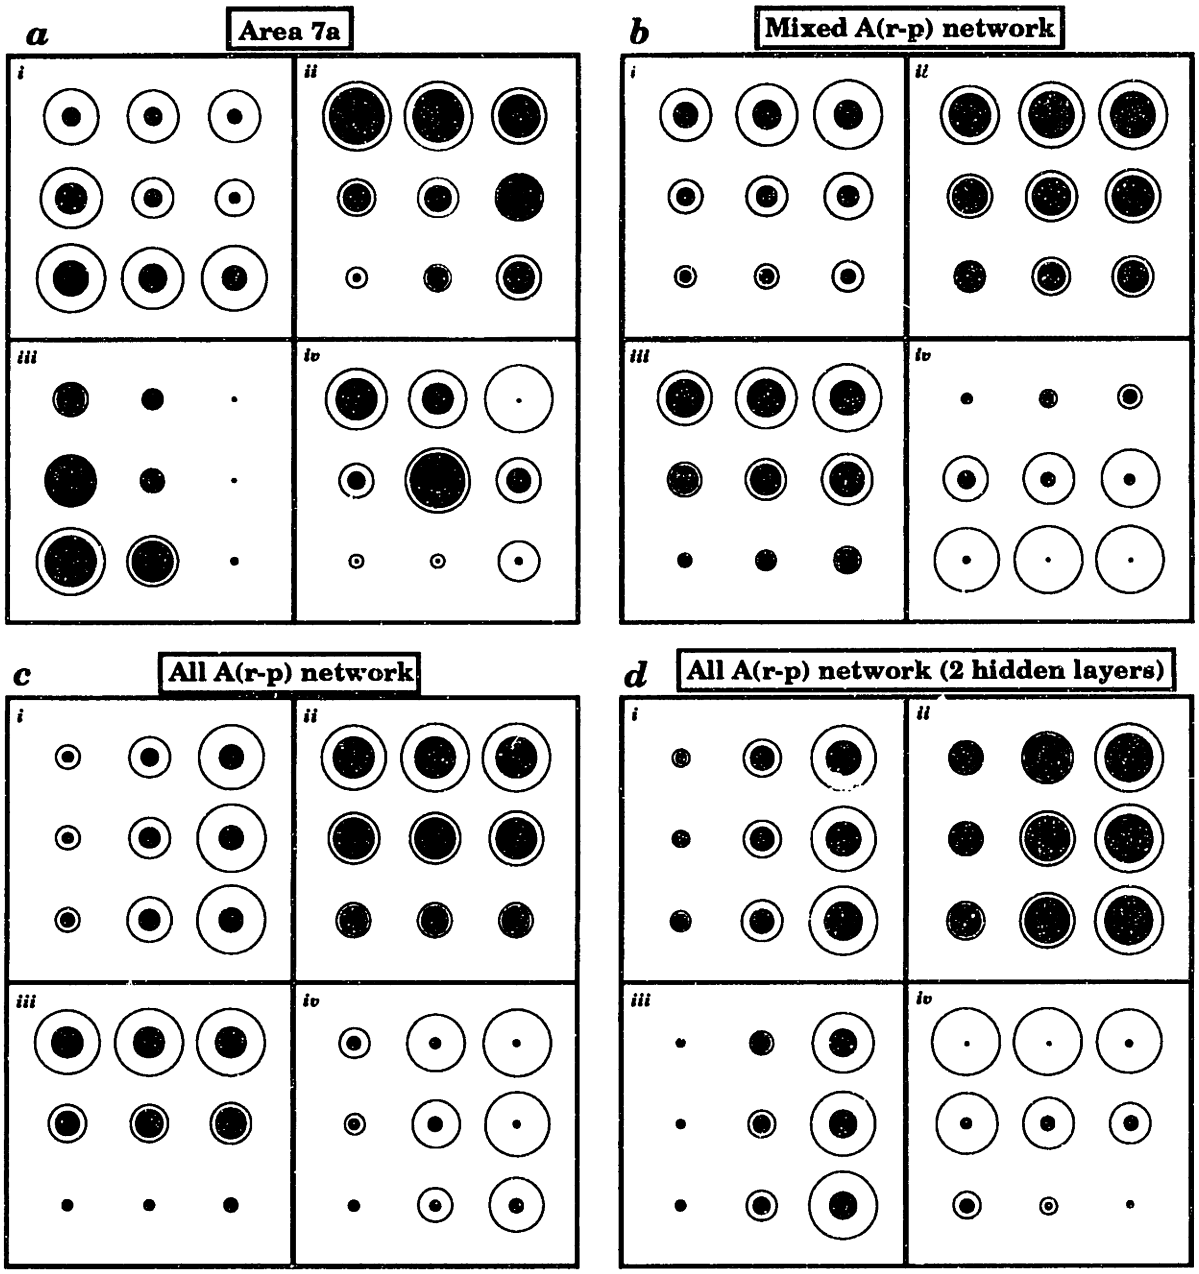
a
Area 7a
i
ii
iii
iv
b
Mixed A(r-p) network
i
ii
iii
iv
c
All A(r-p) network
i
ii
iii
iv
d
All A(r-p) network (2 hidden layers)
i
ii
iii
iv

Figure 7

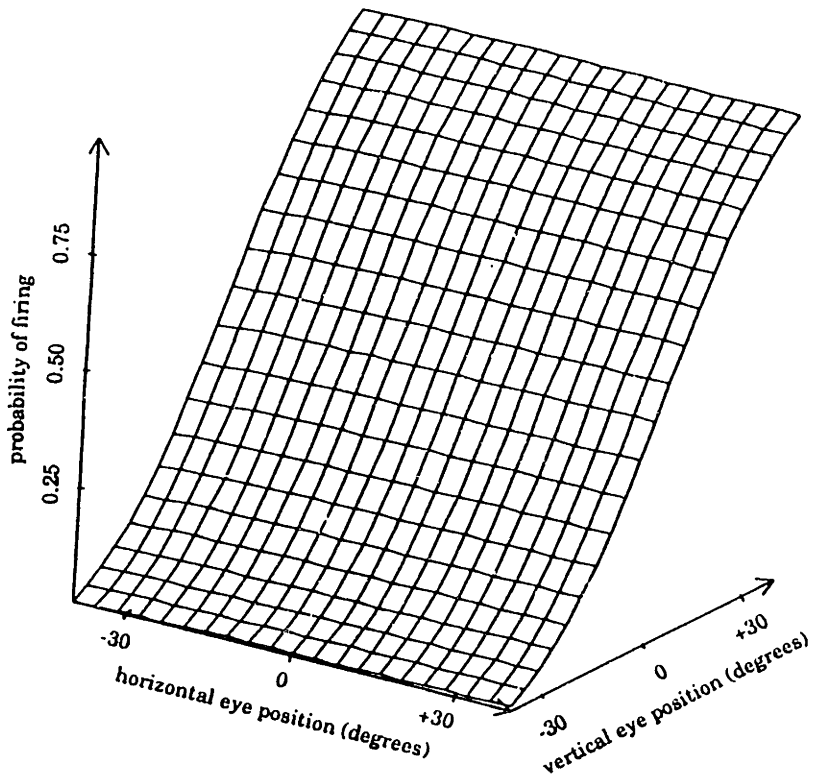
probability of firing
0.25
0.50
0.75
-30
0
+30
horizontal eye position (degrees)
-30
0
+30
vertical eye position (degrees)

Figure 8

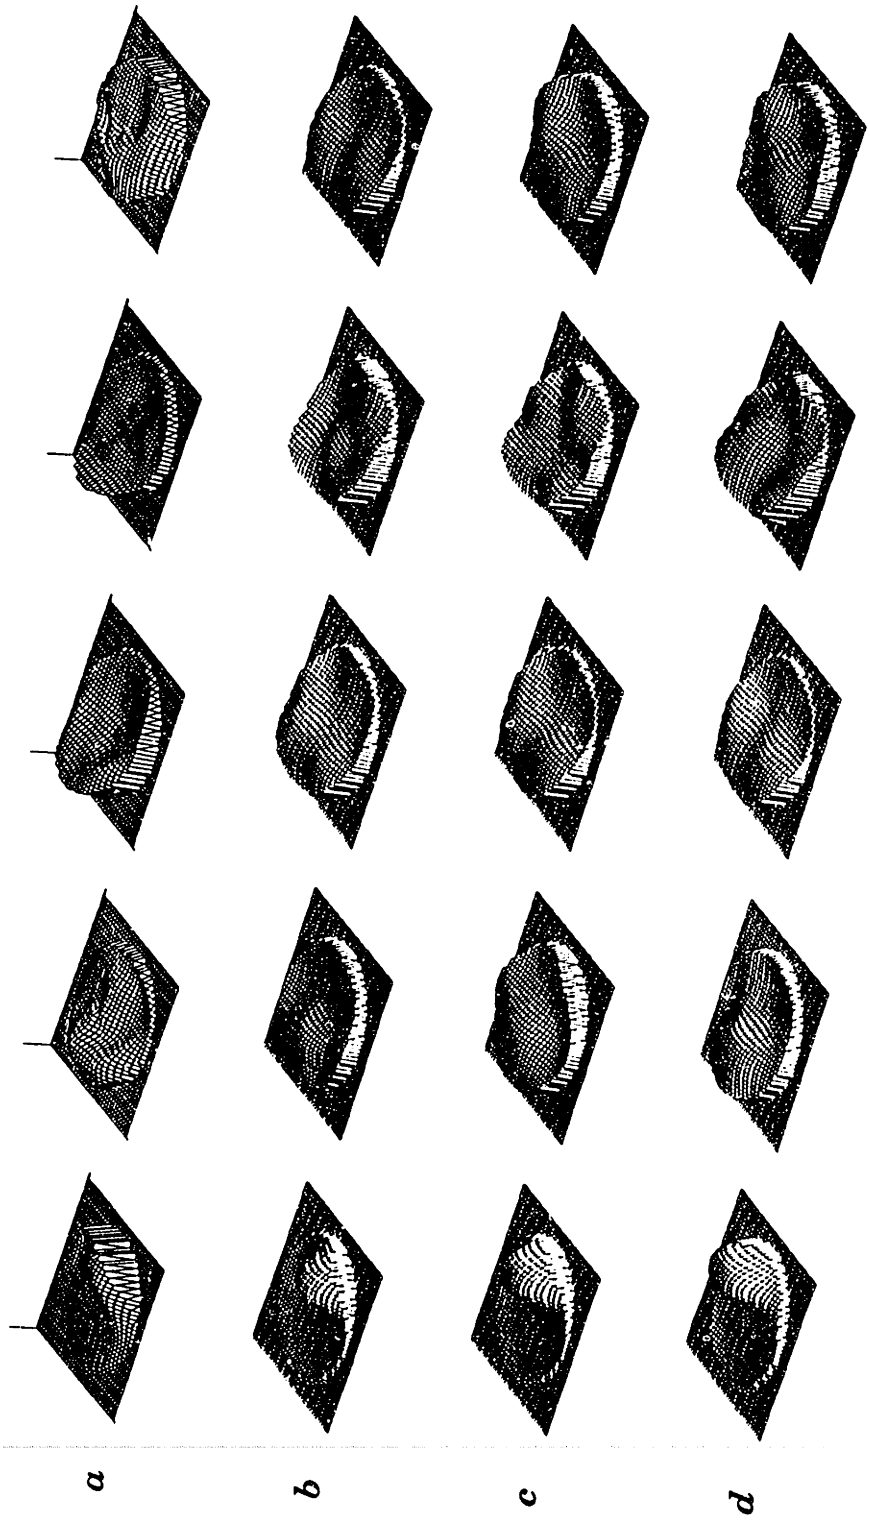

Figure 9

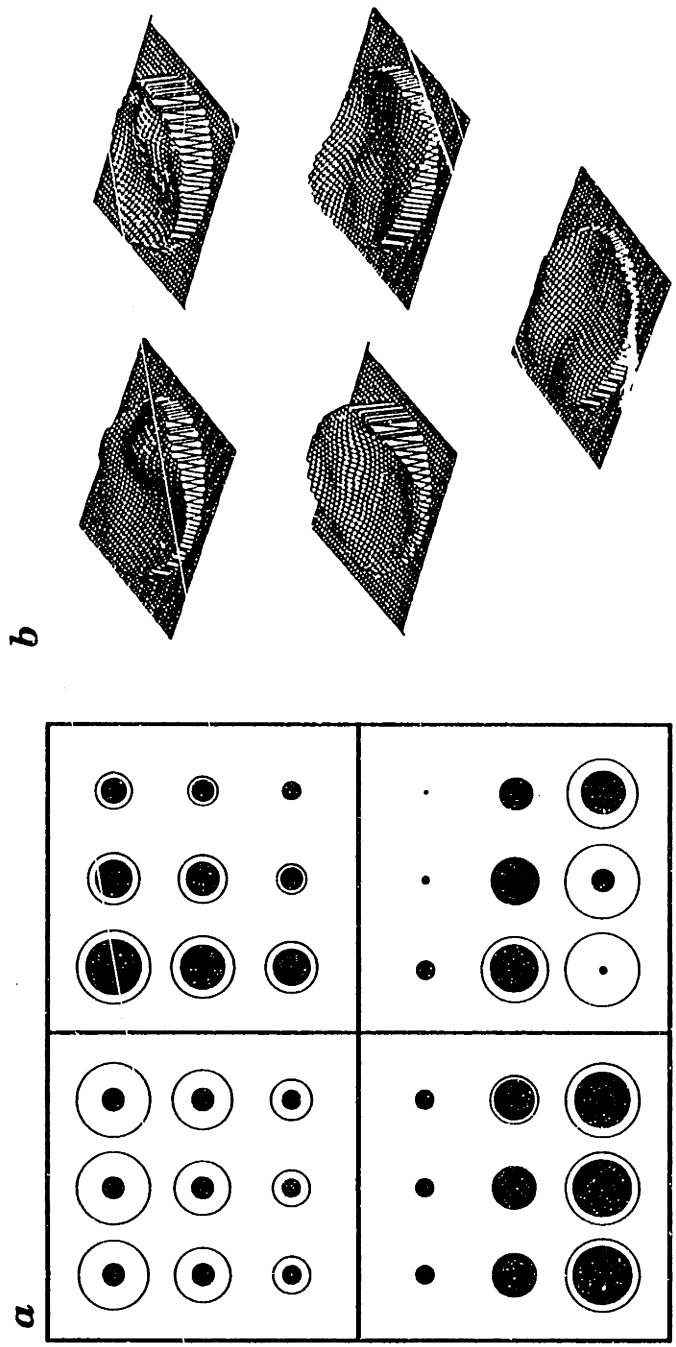

Figure 10

Error (degrees)
0
2
4
6
8
10
12
BP / BP
Mixed Ar-p / BP
Mixed Ar-p / Mixed Ar-p
BP / Mixed Ar-p
BP / BP
All Ar-p / BP
All Ar-p / All Ar-p
BP / All Ar-p
a
b
c
d
e
f
g
h
Continuous training pattern
Binary training pattern

Figure 11

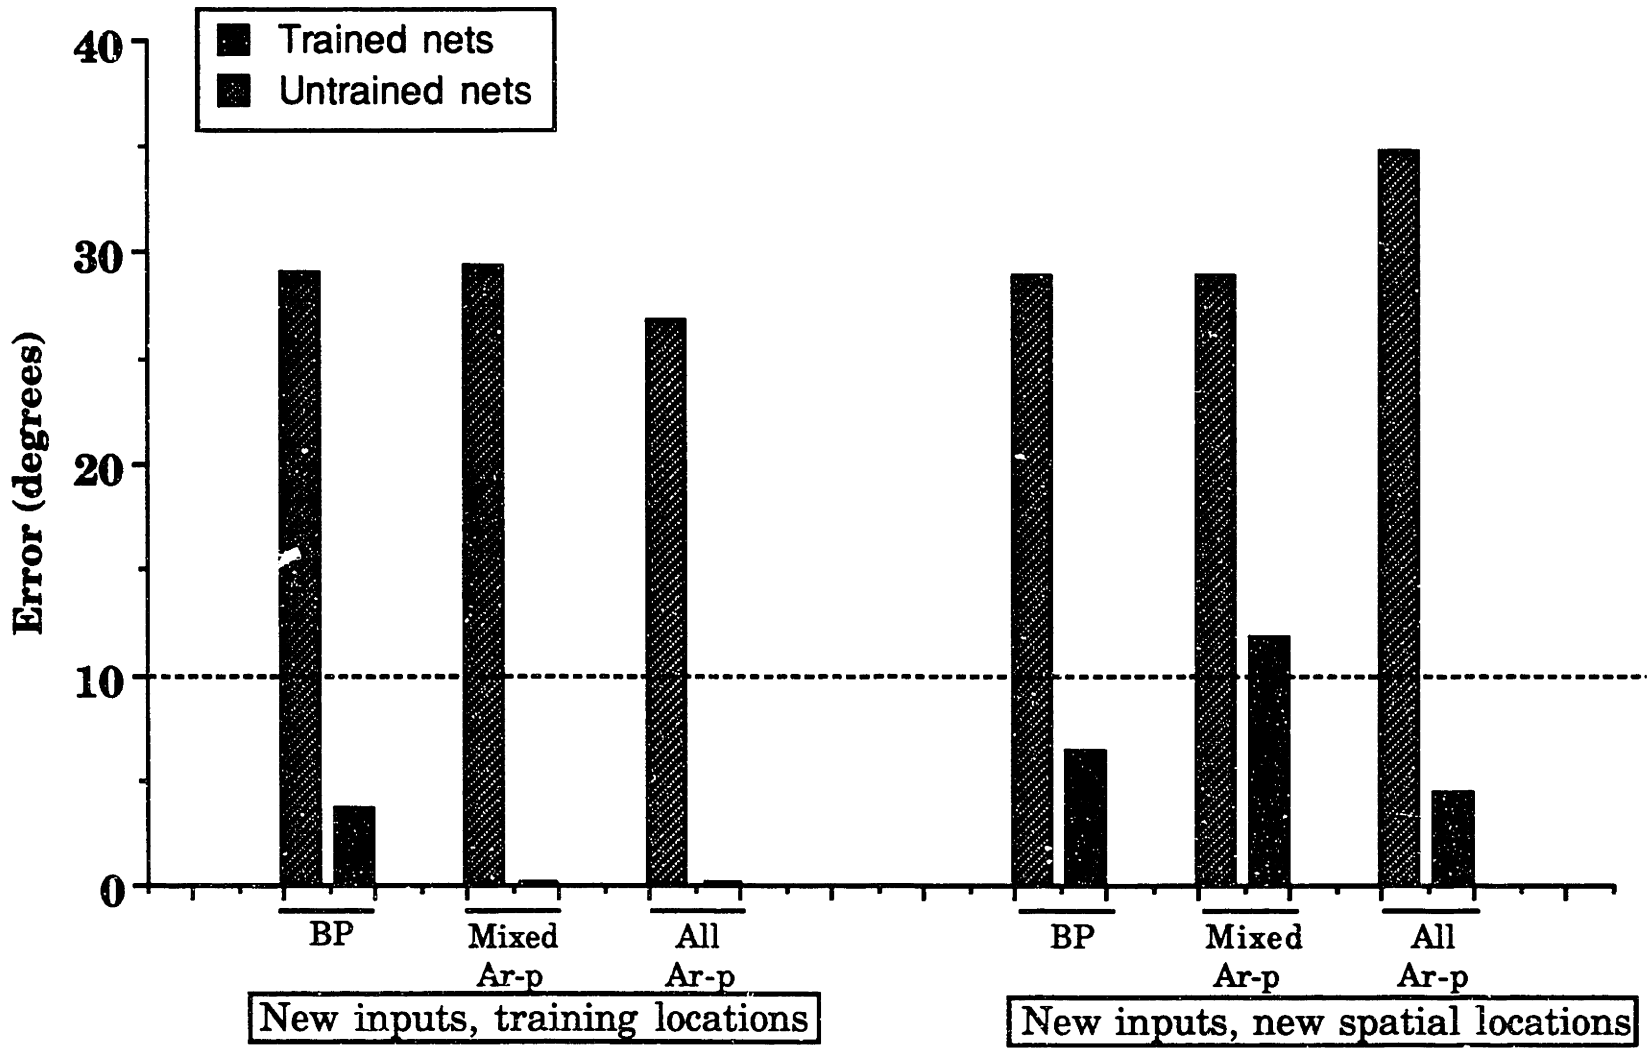
Trained nets
Untrained nets
Error (degrees)
40
30
20
10
0
BP
Mixed Ar-p
All Ar-p
BP
Mixed Ar-p
All Ar-p
New inputs, training locations
New inputs, new spatial locations

Figure 12

II

Integrating Auditory and Visual spatial Information

Chapter 4

Cerebral lesions and the localization of sound sources

SUMMARY

Whereas maps of the visual world are numerous in the primate and the feline visual systems, there are few well-established topographic representations of auditory space. Auditory spatial maps are also generally coarser than those of the visual field, and thus are less obvious candidates for subserving sound localization than their visual counterparts are for the localization of visual stimuli. Studies of the effects of lesions of various parts of the nervous system have thus played a particularly important role in attempting to identify structures that represent auditory space. This chapter is a review of selected studies of the effects of cerebral lesions on sound localization abilities in humans, cats, and monkeys. It will provide a background for our studies—described in following chapters--of the representation of sound location in an area of the monkey's posterior parietal cortex.

I. INTRODUCTION

An important function of the nervous system of vertebrates is to localize the sources of sensory information in extrapersonal space. The visual sense is organized in such a manner that the point of origin of a stimulus is encoded at the very initial stages of processing, by virtue of the focusing function of the lens and of the topographic arrangement of the detectors. As pointed out by Lord Rayleigh (1876), the physical properties of auditory stimuli—specifically, the much greater wavelength of sound waves as compared to light waves—make an analogous arrangement impractical for the auditory system. The ability to localize auditory stimuli, however, may be behaviorally as crucial as visual localization. For example, animals with frontal vision must rely on their hearing to detect the approach of a predator in the entire posterior half of their extrapersonal space. As it turns out, most vertebrates are indeed very good at detecting the spatial origin of sound sources, and they do this by comparing subtle differences in intensity, arrival time, and phase, as well as the spectral composition, of the acoustic signals at the two ears (for reviews, see Blauert, 1983; Yost and Gourevitch, 1987; Syka and Masterton, 1987).

The auditory system of vertebrates is organized so as to allow complex transformations and comparisons of signals from both ears that begin at the very periphery. In fact, binaural information undergoes such extensive processing in subcortical auditory structures that it is not possible to postulate a role of the auditory cortex in auditory localization simply from the arrangement of its inputs. Numerous studies have attempted to elucidate this role using anatomical, behavioral and physiological approaches. I intend to review the neuropsychological portion of this investigative effort. I will describe those studies that have attempted to infer, from lesions in humans and animals, the role of the cerebral hemispheres in general, and of the auditory and association cortex in particular, in auditory localization. In my discussion of the neurological literature I will first focus on earlier efforts to ascribe auditory localization functions to a

specific hemisphere or lobe. Next I will describe cortical ablation experiments in animals that have improved the precision of the functional localization suggested by the early human studies. I will then return to human studies and describe more recent attempts to examine the relationship between these functions and the perceptual representation of spatial relationships.

Three other disciplines that have contributed in an integral manner to our understanding of auditory localization mechanisms are psychophysics, neurophysiology and neuroanatomy. I will not review the results of these studies in detail. I will summarize some of them where appropriate and refer the reader to recent review articles where most of the relevant references can be found.

II. EARLY HUMAN STUDIES: IS THE CEREBRUM NECESSARY FOR SOUND LOCALIZATION?

The study of man' ability to localize sound sources began, unlike many areas of scientific research, under the auspices of a sound physical theory. Given that humans have two ears, it was reasoned at least by the 1800's that a comparison of certain parameters of the sounds arriving at each ear could reveal the direction of that sound's source. The parameters put forth were intensity, phase and arrival time. Lord Rayleigh (1876, 1907) formalized these ideas into a "duplex theory" of auditory localization, which predicted that due to the size of our heads and the velocity of sound in air we might localize high-frequency sounds by comparing interaural intensity differences and low-frequency sounds by comparing interaural time differences (phase and arrival times). This elegant theory stimulated a series of psychophysical experiments (summarized in Sanchez-Longo et al., 1957) which eventually proved it correct as far as behavioral parameters were concerned (Middlebrooks and Green, 1991, for a recent review of human's ability to localize sound). The same theory has also been guiding much of the later neurophysiological research, which in the last two decades has shown that the peripheral and

central auditory pathways do indeed encode the parameters predicted by the theory (see Brugge and Reale, 1985, for a review).

The theoretical background also begged neurologists the question that set the main direction of human studies for the first half of our century: do the cerebral hemispheres play an important role in auditory localization? The fact that certain specific nontrivial transformations need to be performed on the input parameters to compute sound direction suggests that a higher center such as the cortex might be involved. On the other hand, the accurate comparison of parameters with very small differences (e.g., time delays of 0.5 ms), might require the computation to be performed across as few synapses as possible, and thus peripherally.

Unfortunately, the studies that addressed this question in the, 1950's and, 1960's gave mixed results. Walsh (1957) found that 21 patients with various unilateral and bilateral hemispheric lesions had no trouble indicating the apparent direction (left vs. right) of time-delayed clicks presented via headphones. He did find that some patients had impaired ability to localize stimuli separated vertically but no problem in the horizontal discrimination (via headphones). Very few studies of patient populations have since looked, however, at vertical localization capacities. Two problems with this study, pointed out by Shankweiler (1961), were the absence of a control group and the use of a rather long interaural time delay (2.5 ms) for the headphone stimuli, which may have made the task too easy.

Around the same time Sanchez-Longo and his associates introduced a clinical test for sound localization which could easily be implemented in any hospital (Sanchez-Longo et al., 1957; Sanchez-Longo and Forster, 1958). Blindfolded subjects held their heads fixed on a chin rest and pointed to whichever of 13 speakers, arranged radially 15 degrees apart, happened to sound. When compared to a group of normal subjects, 50 patients with focal unilateral cerebral lesions had impaired auditory localization. Most impairments were in the auditory hemifield contralateral to the lesion. Moreover, patients with temporal lesions were most frequently impaired (19/21), although

temporoparietal and frontotemporoparietal lesions also produced impairments (2/3 and 7/7 patients, respectively). The authors concluded that an intact temporal lobe is needed for the localization of sound sources in the contralateral hemispace, and that their test was specific for temporal lobe lesions.

Shankweiler (1961) tried to reproduce Sanchez-Longo's group's results in more controlled conditions and with more patients. He also tested his patients on an auditory discrimination task where they had to point halfway between the perceived directions of two sequentially activated speakers. This test (which required a verbal response) was meant to reveal those patients that might have failed the previous test simply because they had trouble pointing to the sounds. His results were surprising, however. None of the 78 patients showed decreased performance on the localization test, while they were significantly impaired in the discrimination test. Shankweiler deduced that human subjects with unilateral cerebral lesions are not impaired in free-field (i.e., not via headphones) auditory localization.

Several years later the neurologists Klingon and Bontecou took a somewhat drastic approach and applied a very simple bedside test for auditory localization to a series of patients (Klingon and Bontecou, 1964, 1966). The test required the patient to reach with eyes closed to the physician's hand while he rubbed his fingers at some location near the patient. A group of normal subjects and two groups of patients (112 with subtentorial lesions and 42 with supratentorial disease) were tested. The test turned out to be surprisingly sensitive. Auditory localization impairment was found in 33/42 patients with supratentorial lesions, and in neither of the other two groups. Twenty-nine patients in the impaired group had focal unilateral lesions, and their impairment was always contralateral to the lesion. Mislocalization was also more common with right-sided damage, although no comment was made by the authors on this (Bisiach et al., 1984). Klingon and Bontecou thus confirmed Sanchez-Longo's finding that the cerebral hemisphere is necessary for accurate sound localization, and that each hemisphere is responsible for sounds in the

contralateral hemispace. Because mislocalizations were not associated with lesions in any one particular area, however, they rejected the idea that the temporal lobe plays any special role in mislocalization syndromes.

Klingon and Bontecou also rejected Walsh's report that horizontal auditory localization remains intact after cortical lesions, on the argument that auditory lateralization, i.e. distinguishing a left from right sound, is a different task from localization In the latter task the subject must associate a spatial position with the auditory stimulus, and it is this function, they argue, that requires an intact hemisphere. Lateralization can be performed as an exclusion task (i.e., it's enough for a patient to distinguish when a sound is in a field where he/she cannot localize it from a sound that's in the intact hemispace), and thus is likely not as sensitive a test of auditory localization impairments as free-field stimuli are. The same argument has been offered by Jenkins and Merzenich (1984) to explain previous discrepancies in results obtained from humans as well as animals. In fact, Walsh did obtain impairments in vertical localization, which was tested via a free-field stimulus. Sanchez-Longo, who did find impairments, was using a localization test. On the other hand, so was Shankweiler, and he did not obtain mislocalizations.

Another puzzle comes from two case reports by Jerger et al. (1969, 1971). They studied the effects of bilateral temporal lobe damage on two patients who had had 2 successive strokes, one in each hemisphere. Besides experiencing a transient deafness, both patients lost the ability to recognize words or sentences ("word deafness"). While one also lost the ability to localize sound sources, however, the other patient had no trouble localizing free-field stimuli. The authors speculated that an "interaural imbalance in the relation between loudness and signal duration" was responsible for the deficit in the first case. Although this was an interesting idea, to my knowledge it has not been supported by experimental evidence since.

A much more recent study performed by Efron et al. (1983) not only supports a role of the temporal lobe in contralateral auditory

localization, but also suggests a hemispheric division of roles at a more general processing level. These researchers studied the "cocktail party" effect in patients that had undergone unilateral temporal lobectomy. The effect refers to the ability to identify and localize an individual sound among many (and thus be able to choose the most interesting conversation to eavesdrop at a cocktail party). Patients that had undergone temporal lobectomy had to identify, in the first task, all 5 of 5 sounds presented simultaneously and taken from a set of common environmental sounds. The sounds were presented through headphones, and some had intensity differences that allowed subjects to lateralize them. All patients had more difficulty identifying sounds lateralized to the contralateral ear than to the ipsilateral ear. In the second task, the subjects were instructed to search for a specific sound in a set of 5 and to report from which side it was perceived. Again, patients made more errors when the chosen sound was presented contralaterally to the lobectomized hemisphere. These results confirm the temporal lobe's importance in the spatial analysis of the contralateral auditory hemifield. More interestingly, they show that this function may required to perform other types of auditory analysis such as identification. The patients did not show any significant hemispheric asymmetry when the sound where presented monaurally, and hence devoid of spatial characteristics. When a hemisphere was presented with binaural inputs, however, a localization deficit was accompanied by an identification deficit. It was as if any spatial feature of the sound, if present, had to be processed before identification was allowed to proceed; any deficit in localization then seemed to produce poorer identification.

Overall, then, it seems established that an intact contralateral hemisphere is needed for normal auditory localization in each auditory hemifield. The studies I described did not conclusively establish whether the temporal lobe really plays a special role in this function. There remain inconsistencies which hopefully will be resolved by using more refined imaging techniques and more uniform testing procedures.

III. Cortical Ablation Studies In Animals

The earliest studies of the effects of cortical lesions on auditory localization were described by Neff and co-workers (1956), who found that cats with bilateral auditory cortex lesions had great difficulty walking to the source of a sound on a left vs. right choice. In subsequent years, several studies established the effects of bilateral and unilateral cortical lesions on the sound-localizing ability of various species (reviewed in Heffner and Heffner, 1990). Briefly, primates and carnivores are significantly affected by such lesions, while rodents show only minor deficits. In the former two groups, several studies have addressed the nature of the deficits in some detail.

Unilateral lesions

The auditory pathways, unlike the visual system, combine the signals from the two peripheral organs (the cochleae) very early on. Although two parallel pathways starting in each ear follow separate courses all the way to the primary auditory cortex (area AI, in the ectosylvian region of the temporal lobe), there are cross-over connections at several levels in between (see the references cited in Brugge and Reale, 1985, for the original descriptions of the pathways leading up to the auditory cortex). In order to study the contribution of these way-stations to sound localization abilities, Jenkins and Masterton (1982) performed unilateral disruption of inner ear structures, superior olivary complex, lateral lemniscus, inferior colliculus, medial geniculate and auditory cortex of cats. They then tested animals with each lesion on a localization task in which the cat had to walk to the speaker that had emitted a click while the cat was looking straight ahead. There were seven speakers places 30 deg. apart. Lesions below or at the level of the superior olive resulted in bilateral or ipsilateral deficits in sound localization, while all lesions above this level produced deficits only in the sound field contralateral to the lesion. After undergoing ablation of the left auditory cortex, cats performed at chance level for some of the speakers on the right side.

This study thus established that in spite of the non-topographic organization of low-level auditory inputs, these inputs are transformed into aural hemifield representations at the superior olive, and remain segregated—at least in a functional sense—as they ascend to the cortex. The early auditory structures therefore achieve, presumably by analysis of binaural interactions, a topographical segregation of auditory space analogous to the one performed by the optic chiasm for visual inputs. Moreover, the auditory cortex in each hemisphere is necessary for auditory spatial discriminations in the contralateral sound field.

This experiment revealed an interesting aspect of the strategy of cortical-lesioned cats: their performance was worst at a particular point within the contralateral hemifield (at 60 deg. in the example shown by the authors). Such a curve may suggest the presence of a "sigoma," or auditory silent spot, analogous to a visual scotoma. Analysis of the cats' behavior, however, showed that this peculiar performance reflected the animal's strategy to systematically respond to the same sequence of speakers in the impaired hemifield, as this maximized reward in the absence of spatial information. It was important to investigate this issue as the authors did, because obtaining a sigoma from a restricted lesion would strongly argue for the presence of a topographic representation of auditory space.

Jenkins and Masterton also tested their lesioned cats on a simpler, left-right speaker discrimination task, and found no impairment. This result was explained by examining the cats' learning curves in this task and noting that they were simply learning a new strategy (to respond to the contralateral speaker whenever the sound was hard to localize) that did not depend on an ability to localize sounds in the contralateral hemifield. As the authors pointed out, this observation is likely to explain the negative results of previous studies of unilateral lesions, which used left-right speaker discrimination as a measure of auditory localization ability.

Bilateral lesions

Bilateral ablations of auditory cortex produce disruptions in the ability to localize sounds that are more dramatic than those obtained with the unilateral lesions just described. The deficits are permanent and extend to the entire auditory space (e.g., Heffner, 1978; Heffner and Masterton, 1975). Sound-localization mechanisms, on the other hand, seem well-developed within subcortical auditory structures as well. Indeed, cats with auditory cortex lesions can still make reflexive, orienting head movements toward an unexpected sound (e.g., Beitel and Kaas, 1971), and decorticate opossums can indicate the direction of a sound with a non-spatial response (Ravizza and Masterton, 1972). The question then arises as to what is the basis of the cortical deficit. This issue was recently investigated by Heffner and Heffner (1990). They produced bilateral ablations of superior temporal gyri (including all auditory fields) in Japanese macaques (*Macaca fuscata*), and then studied their learning and performance on four versions of a sound localization task. One experimental variable was whether two loudspeakers were located across the animal's midline or within a hemifield of head-centered space. The other variable was the type of response which the animal had to make to indicate which speaker had just emitted a short noise. In the two-choice test, the required response was a spatial one (to walk to the correct speaker); in the "conditioned-avoidance" test, the response was a non-spatial one (to stop drinking if the left speaker had sounded, and to continue drinking otherwise). The monkeys could not discriminate sound locations within a hemifield either in the two-choice task or in the conditioned-avoidance task. In the midline stimulus condition, the animals were much slower than controls in learning to walk to the left or right speaker, and had decreased acuity in the conditioned-avoidance task for these stimuli.

The interpretation given to these results is that deficits in sound-localization abilities produced by auditory cortical lesions derive from both sensory and perceptual deficits. The total inability to

discriminate sound locus within a hemifield and the reduced acuity in left-right discriminations indicate a sensory deficit: the only physical difference between the sounds (their spatial separation) is not detected as well or at all. The difficulty that lesioned monkeys have in learning to approach the source of a sound, as opposed to just indicating its position with a non-spatial response, may indicate an additional perceptual deficit. The sensory system must be detecting the two sound sources as being different, because the monkey can indeed indicate whether the left one sounded by interrupting his drinking. An inability to associate the difference between the stimuli with a difference in spatial locations would explain the difficulty in walking to the appropriate speaker. Such a perceptual deficit might indicate a specific function of the auditory cortex in the processing of stimulus locations: to integrate sensory localization information into some perceptual map of auditory space, or to organize sensory inputs into a format that would allow higher cortical centers to perform the perceptual integration. This integration would presumably not be needed to detect the differences between the stimuli and express this detection in a non-spatial response. Such an interpretation is consistent with the previous findings in other species mentioned above, and replaces the auditory-motor hypothesis of Ravizza and Masterton (1972) (who postulated a general deficit of motor responses to auditory stimuli), which would not explain some of the findings just described.

Partial lesions

Given that the auditory cortex is necessary for contralateral sound localization, how is auditory space represented in this region? Recording studies aimed at elucidating the functional architecture of auditory cortex have failed to uncover a topographic organization like that of the primary visual cortex. The primary auditory cortex (A1) of the cat is divided into rostrocaudally oriented strips of cells with similar responses to binaural stimulation (binaural interaction bands), and dorsoventral/mediolateral isofrequency bands (see Brugge and Reale, 1985, for a review). Early studies had demonstrated neurons

that responded preferentially to sounds presented in limited regions of the sound field (e.g., Evans, 1968), and it was suggested (Imig and Adrian, 1977) that neural sensitivity to spatial location might be organized along the isofrequency bands. A study by Middlebrooks and Pettigrew (1981) examined this possibility by systematically recording along the mediolateral axis of A1 while varying the location of free-field pure-tone stimuli. About half the neurons found in this study were selective for sound location and responded to tones anywhere in the contralateral field (hemifield units) or along the axis of the contralateral pinna (axial units). These two classes of neurons were segregated from each other and separated by neurons insensitive to sound location (omnidirectional units). Although a possible organizational scheme of location sensitivity in A1 was found, it was not a topographic one, and it left open the question of how a population of A1 neurons might encode sound location in a continuous fashion.

In a more recent study Rajan et al. (1990) examined the isofrequency bands of cat A1 in analogous manner, and came up with a more refined classification of azimuth-sensitive neurons. These are indeed clustered according to the region of auditory space to which they are sensitive, and their organization is consistent with a model of a modular arrangement that allows the localization of sounds of a particular frequency by groups of cells within an isofrequency band.

A set of elegant lesion experiments provides further support for such a model of A1 organization Jenkins and Merzenich, 1984). The question asked was, can individual frequency band strips in A1 localize sounds of that particular frequency throughout the contralateral auditory space? First, small lesions limited to a restricted band of frequency in A1 were made. These produced a profound localization deficit for contralaterally presented tones of the corresponding frequency, while localization of all other frequencies remained intact. Then the complementary lesion was made in another cat, i.e., all A1 was destroyed except for a narrow band of frequency. This cat lost the ability to localize contralateral sounds of all frequencies except those in the spared band. This study thus confirmed that A1 in each

hemisphere contributes to only contralateral sound field representation, that it is necessary for normal binaural sound-localization behavior, and established that sound-location representation is indeed organized by frequency channel.

IV. Beyond The Auditory Cortex: Spatial Representations

The evidence from human and animal lesions reviewed so far implicates the auditory cortex in spatial processing of the contralateral hemifield in the auditory domain. How does this auditory information contribute to our perception of external space and to our interaction with the external world? Research on this topic in the various disciplines has been on a variety of schedules.

Psychophysical studies of the interactions among various sensory modalities date back to the early part of this century. There is a vast literature on the myriad combinations of effects that different senses can exert on one another with regards to spatial, temporal and response aspects of sensation. The review by Welch and Warren (1986) is an excellent introduction to this literature. A great deal of evidence underscores a strong interaction between auditory and visual spatial representations. Work by Platt and Warren (1972) and by Lackner (1973), for example, has shown that visual input can affect the absolute calibration of spatial auditory maps, while Knudsen (1985) has shown in the owl that the influence of vision is actually necessary for the proper development of auditory maps. The question remains open, of course, as to whether these two maps coexist separately or combine into one supramodal spatial representation (see, for example, Auerbach and Sperling, 1974).

Physiological studies of central neural processes underlying these issues are still in their infancy, relatively speaking. Jay and Sparks (1984) have demonstrated a gating of auditory receptive fields which aligns these fields with visual maps in the primate superior colliculus. Very few studies, however, have yet looked at the interactions of

sensory maps in more central neural structures such as the association cortices.

Lesion studies got an earlier start on this issue. It seems that around, 1970 neuropsychological research in auditory localization took a turn towards more sophisticated aspects of this function. In particular, several studies have addressed the topic of perceptual representation of auditory space, and its interaction with spatial maps arising from other sensory modalities.

Extinction and neglect

An early report of a "higher-order" auditory spatial dysfunction was a paper by Bender and Diamond (1965), who reported auditory extinction in "more than 25 patients with hemispatial disorientation." Unfortunately their report was very sketchy, and very little information was given regarding the types of patients examined and the tests employed. These authors were also the first to describe "alloacusis" (Diamond and Bender, 1965), a displacement of the acoustic image across the midline, in patients with auditory neglect. In other words, a neglect syndrome was described for the auditory modality, which differs from visual neglect in that the stimuli on the side contralateral to the lesion do not go unperceived, but rather are perceived across the midline on the ipsilateral side.

Auditory extinction was further investigated by Sparks et al. (1970). Right-handed patients with right or left hemisphere damage were presented simultaneous stimuli in the two ears. Right brain damaged (RBD) patients only failed to report stimuli presented to the contralateral ear, while left brain damaged (LBD) patients were divided into a group (18/26) that failed to report contralateral stimuli and a smaller group (7/26) that showed extinction for stimuli presented in the ipsilateral ear. A model was thus proposed, according to which not only does each hemisphere monitor inputs from the contralateral hemifield, but a third pathway connects the left ear with the left hemisphere via the right hemisphere and anterior commissure.

Extinction has been considered an interesting phenomenon mainly because of its association with parietal lobe syndromes involving neglect. This association may not be a true one, however, according to a study by De Renzi et al. (1984) They addressed the issue of whether an attentional deficit or simply a sensory one underlies extinction. They tested 2 groups of patients (RBD and LBD) for auditory extinction, visual extinction, and visual neglect. If extinction were really a sign of an attentional disorder, then the prediction was that it should be supramodal (i.e. visual and auditory extinction should usually coexist); it should be observed in every instance of neglect; and it should show an asymmetry (occur more frequently with right hemisphere damage), as neglect does. None of these features was observed. Auditory extinction was frequently seen independently of visual extinction, and vice versa. No right-left asymmetry was observed. The presence of auditory extinction was indeed found to correlate positively with the presence of visual neglect, but the correlation was not perfect. On the basis of these results, De Renzi et al suggested a sensory deficit as the underlying nature of extinction, and warned against using extinction as a reliable sign of attentional impairments.

While auditory extinction may be a sign of a sensory deficit, auditory neglect is likely to be very closely related to attentional processes and/or distorted spatial representations. Why then not look at the posterior parietal cortex, which is known to be involved in the analogous processes for vision? Two studies by Heilman and his associates (Heilman et al., 1971, 1972) address this issue. In the first study the caudal part of the inferior parietal lobule was ablated in 4 rhesus monkeys. The lesion produced auditory as well as visual and somesthetic neglect syndromes, from which all monkeys recovered within a month. These monkeys showed alloacusis, i.e. they turned to the ipsilateral side when a sound was made on the contralateral side. In the second study 17 patients with auditory neglect were identified by examination. Of 10 of these who had a brain scan, 9 had lesions around the right inferior parietal lobule (1 had a frontal lobe lesion).

They also had other manifestations of neglect (visual, tactile and autonomic). These results directly implicated the IPL in the processing of auditory stimuli (either by spatial analysis or auditory attention); alternatively, they argue for a supramodal attentional and/or representational system subserved by the parietal cortex.

Attention and spatial distortions

The syndrome of contralateral visual neglect associated with posterior parietal cortex damage has been interpreted as part of an impairment of the internal representation of egocentric space (Bisiach et al., 1981). One might then ask how auditory neglect fits in such a scheme. In other words, is auditory spatial perception integrated into a supramodal representation which might subserve mechanisms such as attention, orienting, and navigation? The evidence from cortical lesions so far is mixed.

A study by Altman et al. (1979) first demonstrated a connection between attention disorders and spatial misrepresentations in the auditory domain. By testing auditory localization in depressed patients undergoing electroconvulsive therapy they were able to study the effects of temporary hemispheric inactivation. They studied 96 patients immediately after shocks we applied to one or the other hemisphere. In lateralization tests, where one of two experimenters, each on one side of the bed, addressed the patient, they found that 20 percent of the subjects turned their gaze toward the experimenter on the same side that had been shocked no matter who had spoken. This deficit was similar to the alloacusis phenomenon described by Bender and Diamond, and lasted 1-2 min. It was observed only with inactivation of the right hemisphere. The patients were also tested on lateralization of fused acoustic images (FI) produced by presenting click trains with relative time delays via headphones. While left hemisphere inactivation produced no effect on lateralization, right hemisphere shock caused a shift to the right of the perceived location of all stimuli, including those in the ipsilesional hemifield. For example, stimuli that were normally perceived at -80, -70 and -45 degrees (the

minus sign indicating the left hemifield), were reported by one patient at about -60, -40, -15 and degrees, respectively; stimuli at 0, +45, 70 and +80 deg. were compressed into the right hemifield, at +55, +72, +75 and +85 deg. Similar results were obtained with moving FI's (produced by continuously varying the relative delay between left and right click trains during the presentation). Patients before shock and after left hemisphere shock perceived the sound as moving from left to midline, midline to left, midline to right, and right to midline. Patients after right hemisphere shock, on the other hand, perceived the proper direction of motion, but with the left endpoint right-shifted to about -60 deg. and the midline point shifted to about +40 deg. Right hemisphere inactivation thus seemed to produce a distortion of perceived auditory space such that the entire space was compressed toward the lesioned side. This distortion was topographic, as it preserved the neighboring relationships of stimuli, and generated a region (from about -60 to -90 deg.) towards which responses were not made. Strictly speaking this deficit is not a neglect syndrome, because stimuli elicit responses from all positions in the field. Because these responses are shifted to the right, however, there is a region which is ignored in the motor or response domain.

The results of Altman et al. are partly supported by a study by Bisiach et al. (1984). They studied lateralization of pure tones in a population of 106 patients who had focal brain damage in the anterior/ posterior right/left hemisphere. A systematic directional error (SDI) to the right was found in the right posterior group. This error varied from 0 deg. (at the +90 and -90 deg. stimulus positions) up to +70 deg., and was obtained in both contralateral and ipsilateral hemifields. The errors were larger for right posterior patients that also exhibited visual neglect, although this trend did not reach significance. Based on CT scans, all lesions in these patients clustered around the inferior parietal lobule.

Like Altman et al.'s results, Bisiach's group's study demonstrated a rightward compression of perceived auditory space in patients with right hemisphere damage. This distortion again extended well into the

ipsilesional hemifield. Unlike the previous study, however, the left endpoint of auditory space was not shifted (errors decreased smoothly to 0 towards the left and right endpoints), so that the angular range into which responses were made did not change.

The two studies just discussed both tested lateralization via headphone stimuli. A study by Ruff et al. (1981) addressed the localization of free-field acoustic stimuli. As in Bisiach's study, it was the right posterior damage group that made significantly larger errors in pointing to one of nine hidden speakers. Errors in one hemifield were not significantly different from those in the ipsilateral one, in agreement with the other two studies. Pinek et al. (1989) also examined parietal patients ability to localize free-field sounds. Their subjects had to point to the speaker that emitted a noise or a verbal stimulus. Results were mixed, as only one of three right parietal patients patient showed a significant right-shift in both hemifields (curiously, verbal stimuli were more shifted than noises). The two other patients showed consistent shifts only in the contralateral hemifield. A surprising result of this study was that three patients with left posterior damage found it practically impossible to localize the stimuli, pointing in almost every possible direction with extremely high error rates. It is not clear why these patients performed so badly in both hemifields. The finding is so much outside the patterns of the previous literature that a more detailed investigation with a larger sample of patients will be needed to make some sense of it.

The free-field set-up allowed Ruff to address another interesting question, that of the distinction between personal and extra-personal space. It is possible that personal space (i.e. approximately that within arms' reach) may be represented separately from the rest of the perceptual world, and thus that deficits in the two representations may exist independently of each other. Ruff et al. tested for this by requiring their subjects to point to the halfway between two sound stimuli previously presented in succession. It was argued that this "auditory bisection" task requires the proper representation of spatial locations relative to each other, without reference to the body. A deficit

was indeed found in this function in the right posterior damage group. The errors on this task correlated very well with those on the simple localization task, however, which suggests that if separate representations do exist, they are co-localized in the right posterior hemisphere.

It is natural at this point to attempt a comparison with visual perception and the effects of brain damage on spatial perception in this modality. These have been reviewed extensively (see, for example, Andersen, 1987; De Renzi, 1982; Hyvärinen, 1982; Critchley, 1953). Two patient studies that have specifically looked at spatial representations after parietal damage support the view that there is a general distortion (in the form of an ipsilesional shift) encompassing both hemifields (Ratcliff and Davies-Jones, 1972, and Corin and Bender, 1972, both cited in Bisiach et al., 1984). Most reports of visual disorientation and neglect syndromes, however, describe only contralateral deficits, and the involvement of an attentional process has been clearly established (e.g., Posner, 1984). The question of the relative contributions of spatial vs. attentional deficits in parietal lobe visual syndromes remains open. It is just the issue of the interaction of attentional processes and spatial representations, however, that has repeatedly raised the question mentioned at the beginning of this section, namely concerning the existence of modality-independent neural representations of external space.

It has been argued that an attentional system, in order to be useful as such and efficiently allocate sensory resources, requires "supramodal," that is, modality-independent, sensory representations. These representations would allow spatial attention, for example, to modulate perceptual processes according to a single spatial map integrated across visual, auditory and somatosensory inputs. The systematic and bilateral distortion of perceptual auditory space revealed by the studies described above (Altman et al., 1979; Bisiach et al., 1984; Ruff et al., 1981), and of visuospatial perception (cited in Bisiach et al., 1984), together with the hemispheric asymmetry of these deficits, would seem to suggest a

shared mechanism of spatial representation on which the visual and auditory systems converge, which is distorted by right hemisphere damage. This hypothesis would predict visual and auditory spatial deficits to co-occur with high correlation after posterior hemispheric damage. The degree of this correlation is not known, but appears to be low. It has been traditionally considered low for two reasons: gross auditory localization deficits were not observed in patients with obvious visuospatial deficits, and visual neglect is often observed without visual neglect (see the study described above by De Renzi et al., 1984). Auditory localization studies such as the ones described above, however, have pointed out that deficits of this function tend in general not to be as apparent as visuospatial ones, and can be missed by simple right/left discrimination tasks.

A study that addressed the relationship between sensory modality and attention was recently performed by Farah et al. (1989). Eight right parietal patients were tested with a variant of Posner's spatial cuing task in which a visual or auditory cue appeared in the correct or incorrect location of an immediately following visual target. Posner's (1984) basic finding had been that parietal patients take longer to move their gaze to a contralateral visual target after an invalid visual cue (i.e., a cue the had appeared on the right). This result had led Posner to characterize the attentional deficit that follows parietal damage as a difficulty in disengaging attention from the ipsilesional visual field. Farah et al. found that the same result generalizes to auditory cues: her subjects showed an increased reaction time to left visual stimuli presented after an invalid visual or auditory cue. The authors interpreted this result to imply that the parietal lobe's attentional mechanism does indeed operate on a supramodal representation of space.

Farah's result is a substantial piece of evidence in favor of modality-independent attentional mechanisms. Its implications for the modality of spatial representations are more debatable. The spatial element in the stimulus she used consisted of a right vs. left discriminations, and, as was pointed out for auditory studies, such

discriminations may require very little use of spatial sensory maps. What is needed next is a systematic study mapping of perceived stimulus locations (like Bisiach et al.'s) for both vision and audition in large population of patients with focal hemispheric lesions.

V. CONCLUSION

I chose to review lesion studies of auditory localization because they seem to have provided, so far, the most direct analysis of the possible roles of cortical structures in auditory localization. Recording studies are yielding a wealth of information about the mechanisms underlying these possible roles, and psychophysical experiments continue to refine our understanding of the physical and perceptual parameters that are relevant for these processes. Lesions, however, allow us to directly examine the functional significance of individual structures for specific abilities. Animal studies have shown that the auditory cortex is a necessary structure for spatial discrimination in the contralateral sound field, and that its functional organization matches that suggested by recording studies. Studies of brain-damaged patients have offered insights into the role of auditory space perception in more abstract perceptual and cognitive processes. Lesions of the posterior parietal cortex do not present a consistent picture, but suggest that auditory space representations may be maintained separate from visual ones at several levels among cortical areas, which may reflect as yet uncovered specialized processing of the auditory world.

Chapter 5

Spatially tuned auditory responses in area LIP of monkeys performing memory saccades to acoustic targets

SUMMARY

The lateral intraparietal area (area LIP) of the macaque's posterior parietal cortex (PPC) lies in the dorsal stream of extrastriate visual areas. It receives extensive visual inputs and sends outputs to several eye movement centers. It contains neurons with visual and saccade-related responses suggesting a role of area LIP in programming saccadic eye movements to visual targets. Because primates can also orient to non-visual stimuli, we asked whether LIP neurons process stimuli of other modalities besides the visual one by comparing their activity in auditory and visual saccade tasks.

We recorded the activity of single neurons of *Macaca mulatta* monkeys while they performed memory saccades to acoustic and visual targets. We analyzed the activity during stimulus presentation (stimulus period, S) and during the delay (memory period, M) between stimulus presentation and the saccade to its remembered location.

Among 80 area LIP neurons tested we found 44 that had S period and/or M period responses following presentation of the auditory stimulus. Most of these responses were selective for the left or right stimulus location (spatially tuned) (27 of 29 S responses; 25 of 29 M responses).

The majority of neurons with responses in the auditory memory saccade task also responded in the visual version of the task. 89% (24/27) were clearly bimodal in the S period and 88% (23/26) were bimodal in the M period.

Almost all the neurons with spatially tuned auditory responses that were bimodal were also spatially tuned in their visual responses (20/22 for S responses; 18/19 for M responses). The spatial tuning was usually the same for the two modalities: in 85% (17/20) of the tested neurons for the S responses, and in 83% (15/18) for the M responses.

Area LIP contains a population of neurons that respond to both visual and auditory stimuli. This result is consistent with our finding that the memory activity of most LIP cells encodes the next planned saccade (Bracewell et al., 1991; Mazzoni et al., 1992; see Chapters 7

and 8). If cells are coding planned movements, they should be active independently of the sensory modality of the target for the movement, as was the case for most of the neurons described in the present study.

INTRODUCTION

The posterior parietal lobe of the primate brain contains cortical regions that are specialized for localizing visual stimuli (see Andersen, 1989 for review). Lesions of this brain region in humans and monkeys produce a range of deficits in the perception of a visual stimulus' location as well as in the programming of orienting and reaching movements toward the stimulus (Balint, 1909; Brain, 1941; Critchley, 1953; De Renzi, 1982; Holmes, 1918; Lynch, 1980, 1989, 1992). For example, one symptom of a bilateral posterior parietal lesion is difficulty in voluntarily shifting gaze from one visual stimulus to another (Balint, 1909).

The response properties of posterior parietal neurons suggest that their activity may be the neural substrate for a variety of sensorimotor integration abilities. Recording experiments have recently pointed to an area within the PPC (area LIP) that may play a special role in programming orienting movements to sensory stimuli. Area LIP was identified as an extrastriate visual area in the occipito-parietal cortical visual pathway, i.e. in the dorsal stream of cortical areas that process mainly the spatial aspects of the visual scene (reviewed in Mishkin et al., 1983). It receives extensive inputs from several visual areas in the occipital, temporal and parietal lobes, and projects to higher-order visual areas in the parietal lobe, to association and premotor cortical areas in the frontal lobe, and to the intermediate and deep layers of the superior colliculus (Lynch et al., 1985; Andersen et al., 1990a, 1992; Blatt et al., 1990). Anatomically this area is a link between the early stages of cortical visual processing and premotor, motor, and cognitive centers.

The response properties of LIP neurons also suggest a role in visuomotor processing. When a monkey is trained to memorize the location of a visual stimulus and to look at that remembered location after a delay (memory saccade), LIP neurons modulate their activity selectively for particular stimulus locations and eye movement directions (Gnadt et al., 1988; Andersen et al., 1990b, 1992; Blatt et al.,

1990; Barash et al., 1991a, b). Various proportions of the population of LIP neurons respond during the appearance of the visual stimulus (stimulus period), during the saccade toward that stimulus' location (movement period), and during the delay between stimulus appearance and saccade (memory period). These responses are spatially tuned, that is, the neurons have visual, memory and motor fields, and the three types of fields generally overlap for a given neuron. Gnadt et al. (1988) and Barash et al. (1991a) hypothesized that these response properties reflect the processing of a sensory stimulus for the programming of a motor plan to orient to that stimulus.

The activity of LIP neurons has so far been recorded in tasks that use visual stimuli as cues to spatial locations. Primates can, however, localize stimuli of at least two other modalities (auditory and somatosensory). We asked therefore whether area LIP is involved in localizing nonvisual stimuli. We recorded the activity of LIP single neurons in monkeys trained to perform memory saccades to auditory targets as well as to visual targets. We found a population of neurons that have spatially tuned activity in the auditory memory saccade task. Most of these neurons responded similarly during the auditory and visual versions of the task.

METHODS

Animals, surgery and animal care

Two adult male *Macaca mulatta* monkeys were used in this study. Through a surgical procedure a metal head-post was mounted in dental acrylic on the monkey's skull and a scleral search coil was implanted in one eye (Judge et al., 1980; Robinson, 1963). The monkeys were trained via operant-reinforcement techniques in several saccade tasks including the ones used for this study. In a second surgical procedure a recording chamber was mounted over the posterior parietal cortex (Brodmann's areas 5 and 7). Several months later a second chamber was mounted over the posterior parietal cortex of the other hemisphere. All surgical procedure were carried out under general

anesthesia (10 mg/kg intramuscular ketamine followed by 10 mg/kg intravenous pentobarbital sodium (Nembutal) using aseptic techniques. After each procedure the monkeys received analgesics and systemic antibiotics and rested for a week.

During the training and recording periods the monkeys were deprived of water in their home cages and received apple juice or water as reward for correct task execution, supplemented by additional water at the end of each session to reach the required daily ration. Each monkey had at least two days of rest per week with unrestricted water access. The monkeys received routine veterinarian care. Their well-being was observed in accordance with National Institutes of Health guidelines.

Experimental set-up and behavioral tasks

The monkeys learned to perform several tasks involving saccades for the purposes of several studies. The ones used in this study are the auditory memory saccade task and the visual memory saccade task. The monkey sat in a completely dark room facing a large featureless tangent screen placed 57 cm away. Small light spots (~0.5 deg. diameter, ~45 cd/m^2) could be back-projected onto the screen from two projectors through galvanometer-controlled mirrors. A trial started when a spot was turned on directly in front of the monkey and the monkey started fixating on it. After 750 ms a stimulus was then presented for 750 ms to the left or to the right of the fixation spot at an eccentricity of 8 degrees. The monkey had to continue looking at the fixation spot for another 1250 ms after stimulus offset. At this point the fixation spot was extinguished and the monkey had to make a saccade, in the dark, to the remembered location where the stimulus had appeared. In the auditory memory saccade task the stimulus was a 20-20,000 Hz white-noise burst (70-80 dB sound pressure level) from one of two speakers, each located in front of the tangent screen and 10 deg. to the right or to the left of the fixation point. In the visual memory saccade task the stimulus was a second light spot that appeared 8 deg. to the left or to the right of the fixation point. We

pseudorandomly interleaved left/right and auditory/visual stimulus presentations. A laboratory computer (DEC PDP 11/73) presented the stimuli and monitored the monkey's behavior.

Data collection and analysis

We recorded eye position with the scleral search coil method (Judge et al., 1980; Robinson, 1963), sampling at 500 Hz. We recorded the extracellular potential of single cortical neurons with glass-coated Pt-Ir microelectrodes (Wolshbart et al., 1960) mounted on a Chubbuck microdrive. The laboratory computer stored the eye position samples and the time of occurrence of action potentials for off-line analysis.

We analyzed quantitatively neural activity during three periods of the task defined as follows: the background (B) period was the 450 ms before stimulus appearance while the monkey fixated straight ahead; the stimulus (S) period was from 100 ms after stimulus onset until stimulus offset; the memory period (M) was from 100 ms after stimulus offset until the fixation point's offset. If a cell's response clearly started well within one of these periods we redefined the period so that it consisted mostly of the cell's response time. For each cell we computed a background firing rate B_0 equal to the average firing rate in the background periods preceding left/right auditory/visual stimuli. We defined a response as a significant change in the average firing rate during the S or M periods relative to B_0 (two-tailed t test, alpha level 0.05). If a cell had a response in the S or M periods, we considered the response spatially tuned if the response followed only right or left stimulus presentation, or if the response in a left trial was significantly different from the response in a right trial. For cells with clear S-period responses we estimated the response latency from plots of the firing rate histogram made with 20 ms bins.

Histology

The neurons described in this study were isolated in area LIP of the right hemispheres of two monkeys. After several months of recording the monkeys were killed and the neurons' locations were reconstructed

based on each penetrations' chamber coordinates and depth relative to various landmarks. These landmarks consisted of several DC electrolytic lesions made in the last few weeks before sacrifice, fluorescent dye injections, and guide wires inserted in the brain before sectioning it (for details see Barash et al., 1991a).

RESULTS

Database

We recorded the activity of 80 neurons in area LIP of two hemispheres of two monkeys while they performed the auditory and visual memory saccade tasks. Among these neurons we found 44 that had significant responses in the stimulus and/or memory periods of the auditory memory saccade task. These form the database of our study.

Stimulus-period responses

We found auditory S responses in 66% (29/44) of the neurons in our database. Fig. 1a shows such a response to sound from the right speaker. The response is absent when the sound comes from the left speaker (Fig. 1b), i.e. the response is spatially tuned. The auditory S responses were spatially tuned in 93% (27/29) of these cells.

Most of the cells with auditory S responses that were tested on the visual task also responded to the visual stimulus (89%, 24/27; Fig. 1c). In some cells the visual response was larger while in others the auditory response was larger. The onset and time course of the visual and auditory S responses was very similar in some cells and rather different in others (see *Latencies* below). The visual S responses were spatially tuned in 88% of these cells (21/24; Fig. 1c vs. Fig. 1d), and the spatial preference for the visual and auditory stimulus was most frequently the same (85%, 17/20 cells).

The cells without significant visual S responses were not purely auditory: they did have small responses to the visual stimulus that did not reach significance. These responses had the same spatial preference as the auditory ones.

Memory-period responses

We found auditory M responses in 66% (29/44) of the neurons in our database. Figs. 1a and 2a show M activity after presentation of the sound from the right speaker while the monkey plans a rightward saccade. This activity is absent between the left sound presentation and a leftward saccade (Figs. 1b, 2b). The auditory M responses were spatially tuned in 86% (25/29) of the cells.

Most of the cells with auditory M responses that were tested on the visual task also had visual M activity (88%, 23/26; Fig. 2c). As with the S responses, some cells had more M activity in the visual task while others had more M activity in the auditory task. The time course of the M activity in the auditory and visual tasks was very similar in some cells and rather different in others. In almost all cells, however, the activity increased monotonically (or decreased if inhibitory) and remained steady until the fixation point was extinguished. The visual M responses were spatially tuned in 91% of these cells (21/23; Fig. 2c vs. Fig. 2d), and the spatial preference for the visual and auditory stimulus was usually the same (83%, 15/18 cells).

Of the cells without significant visual M responses, two had some visual M activity that did not reach significance, while one had no visual M activity.

Coincidence of auditory S and M responses

Of the 44 cells with auditory S and/or M responses, 41% (18/44) had both S and M, 32% (14/44) had only S, and 27% (12/44) had only M. Compared with the analogous proportions for visual responses reported by Barash et al. (1991a), we found more cells with only M responses (13% in their study) and fewer cells with both S and M responses (58% in their study).

Fig. 3 summarizes the distributions within our database of the responses properties described so far.

Latencies of the S responses

For a set of neurons with clear S responses we measured the latency of S response onset relative to stimulus onset. Latencies of auditory S responses ranged from 30 to 250 ms with a median value of 155 ms (Fig. 4a), while visual S response latencies ranged from 60 ms to 210 ms with a median value of 125 ms (Fig. 4b). The auditory and visual latencies are thus rather similar across the neuronal population. These values are also similar to the latencies of visual S responses of LIP neurons reported by Barash et al. (1991a).

We also compared auditory and visual latencies within individual cells that had clear S responses to both auditory and visual stimuli. There was no systematic pattern in the differences of latencies, some cells having earlier auditory responses and other cells having earlier visual responses (Fig. 4c). The distribution ranged from -90 to, 190 ms with median 0 ms. An important note of caution when comparing the auditory and visual S response latencies is that we did not attempt to match the perceptual saliency of the sound and light stimuli (see the *Discussion*).

Saccade-related activity

The majority of LIP neurons we isolated had responses related to the saccade, as was observed by Barash et al. (1991a). Fig. 5 shows an example of this activity. Of the neurons with tuned auditory and visual memory responses 73% (11/15) were active during the saccade, the activity always beginning before the saccade. The saccade-related responses were largely similar for the auditory and visual saccades. Moreover, when a neuron had saccade-related responses, these were present in both the visual and auditory memory saccade tasks.

Other response properties

The neurons in this study had excitatory as well as inhibitory responses in the auditory and visual S and M periods. About two thirds of the responses of all types were excitatory (18/28 auditory S,

19/25 visual S, 20/26 auditory M, and 18/28 visual M), and the remainder were inhibitory.

Among the neurons with spatially tuned responses, both contralateral and ipsilateral preferences were represented, but contralateral preferences were more common. Of the excitatory responses, about two thirds of the S period responses were contralateral (12/18 auditory S, 12/19 visual S), while just over half of the M responses were contralateral (11/20 auditory M, 10/18 visual M). We report the proportions for excitatory responses for comparison with the percentages reported by Barash et al. (1991b). These are slightly larger for the S responses (71% in their study) and larger for the M responses (69% in their study). The proportions of contralateral and ipsilateral preferences changed only slightly when we included the inhibitory responses.

DISCUSSION

Our main findings are that, (i) a population of LIP neurons respond during the stimulus and memory periods of memory saccades to acoustic targets; (ii) most of these neurons respond similarly in the visual memory saccade task; and (iii) the responses of most neurons are spatially tuned, with matching spatial tuning for the auditory and visual modalities.

Stimulus-period responses

The S responses are very likely sensory responses to a sound in the neurons' auditory receptive fields. We believe so because they usually start within 200 ms of sound onset, they often have a sharp onset when aligned with sound onset and, in cells without memory activity, they wane before or soon after stimulus offset. They are also almost always accompanied by S responses to visual stimuli. The latter persist in tasks that require no behavior except fixation (Andersen et al., 1985b, 1987; Mountcastle et al., 1975; Robinson et al., 1981), and have thus been considered sensory responses.

These sensory responses could be auditory responses, i.e. arise specifically from the presence of an acoustic stimulus, or supramodal responses, reflecting the presence of a sensory stimulus independently of its modality. The fact that most cells respond to visual as well as auditory stimuli argues for the latter interpretation. Nonetheless we consider the auditory and visual S responses distinct, mainly because they have different time courses in most neurons and thus are likely to arise from separate input channels.

It is important in this regard to note that we did not attempt to equalize the perceptual saliency of the auditory and visual stimuli in our experiment. A difference in saliency could produce responses with different time courses or response magnitudes. In particular, the latency of responses to visual stimuli of retinal, LGN and area V1 neurons depends on the stimulus' brightness. Because we never changed the intensity of our speakers' sound or the brightness of our visual stimuli, however, we would expect a difference in stimulus saliency to produce a *constant* difference in the time course of the response—e.g., all responses to sounds to have longer latencies than the responses to lights. Instead we found some cells with shorter auditory latencies and other cells with shorter visual latencies.

A difference in the time course of auditory and visual responses could also be due to a difference in the process of spatial localization of auditory and visual stimuli. Thus the auditory S responses might build up more slowly than the visual ones simply because it takes longer for the monkey to localize an acoustic stimulus. If this were the case, however, we would again predict a relatively constant difference between the time courses of S responses to sounds and lights across all cells. We instead observed some cells with brisker responses to sounds and other cells with earlier and sharper responses to lights. These cells thus have distinct responsiveness to auditory and visual stimuli.

A more remote possibility is that LIP neurons are purely visual and that their S responses to acoustic stimuli reflect not the auditory stimulus but a visual image of its location. The monkey could in principle localize the speaker's location by imagining its location in a

mental visual map, and cells with visual sensitivity could become active as if a visual stimulus had appeared. This would be a very interesting result because it would support certain hypothesized mechanisms of mental imagery (Kosslyn, 1988) and because no one has yet reported activity of single visual neurons during mental imagery. Without further evidence addressing this hypothesis directly, however, we prefer the simpler interpretation that the S responses we observed are elicited by the physical stimulus present (the sound) and not by a visual image of it; that is, we consider them auditory responses.

The fact that the neurons we studied have distinct auditory and visual responsiveness does not mean that their purpose is to maintain separate auditory and visual codes. The different time course of responses in each modality indicates distinct sources. Other neurons reading these responses, however, may discard differences in their time course and use only the average firing rate of a population of LIP neurons as information. The neurons with auditory and visual S responses may thus integrate auditory and visual information into a code that effectively supramodal and indicates the presence or absence of a localized stimulus. This code could be used, for example, to generate the M-period activity which would in turn be used to program a saccade to the stimulus, as described below.

We have referred to the auditory responses we observed as spatially tuned because they were different for sounds coming from speakers at different locations. An alternative interpretation might be that the neurons are tuned for sound frequency rather than sound location, and that the difference in responses reflects difference in the frequency spectra of the sounds produced by the two speakers. Frequency tuning would not be too surprising given that the majority of neurons with auditory responses in the central nervous system do depend on stimulus frequency. We cannot exclude this possibility but we consider it unlikely. The noise bursts which we used as auditory stimuli contained a very wide spectrum of frequencies, so that each frequency constituted only a small component of the power spectrum. If a neuron were responding to a particular frequency band, the difference in

intensity at this band between the two speakers would thus also be rather small and unlikely to elicit differences in responses as large as the ones we observed. In fact, the noise produced by the two speakers did not sound obviously different. We could not discriminate between the two speakers based on their sound quality while sitting in the darkened set-up. Moreover, the speaker preference for almost all the bimodal neurons matched their visual receptive field. This result would be quite a coincidence if the speaker preference were due to frequency tuning. We consider it more likely that these neurons' speaker preference reflects tuning for sound source location. A similar argument can be made to exclude intensity tuning as the source of speaker preference.

In summary we interpret the S responses as auditory and visual sensory responses that encode the location of auditory and visual stimuli.

Memory-period responses

A group of cells in our study had significant activity during the M period of the auditory and visual versions of the memory saccade task. We refer to these as auditory and visual memory responses, respectively, because in spite of often being very similar they can still be distinguished within each cell by their time course. As for the S responses, we maintain this distinction to indicate that the two activities likely arise from distinct input processes. Also as discussed with regard to the S responses, this still does not exclude that the auditory and visual M responses may play the same role for the two modalities by transmitting the same information (e.g. their average firing rate alone) about the visual and auditory stimuli.

The M period activity could a priori reflect a number of processes. During the delay the monkey must maintain fixation, remember the stimulus' location, shift his attention to that location, and plan a saccade of the appropriate size and direction. The memory activity cannot reflect the maintenance of fixation because it is in general

different while the monkey, always fixating straight ahead, remembers and plans saccades to stimuli at different locations.

The hypothesis that the memory activity reflects a shift of visual attention is an interesting one because shifts of attention can indeed modulate the visual responses of PPC neurons (Bushnell et al., 1981) and because attentional deficits are a prominent symptom of parietal lobe damage. The syndrome of unilateral neglect, of which the inability to shift attention while fixating is at least one component (Posner et al., 1984), can include inattention or decreased attention to auditory stimuli in humans (Heilman and Valenstein, 1972) and monkeys (Heilman et al., 1971; De Renzi et al., 1984). It is unlikely, however, that the memory activity we observed reflects an attentional shift. We have shown in a separate study that LIP neurons' memory activity does not appear if an attended location, cued by a visual stimulus, is not the target of the next saccade (Bracewell et al., 1991; Mazzoni et al., 1992).

The remaining two hypotheses cannot be distinguished by the results of this study. The memory activity could reflect the monkey's memory of where the stimulus appeared, or it could reflect the plan for the next saccade to be made. In the first instance our results would show that the neurons store bimodal memory traces, or an abstract supramodal memory of a spatial location. If the activity reflects the next saccade's plan, on the other hand, then our findings support the ability of these neurons to generate saccade programs irrespective of stimulus modality. In a separate study in which monkeys made two consecutive memory saccades to locations defined by visual targets (Bracewell et al., 1991; Mazzoni et al. 92) we have shown that a component of area LIP memory activity reflects the next planned saccade. We thus favor the hypothesis that the auditory and visual memory activity we observed in this study reflects the covert process of programming a saccade to a particular location, regardless of how that location is specified.

Saccade-related responses

A detailed analysis of the saccade-related activity of the neurons in our database was beyond the scope of this study. One relevant result is that many neurons had saccade-related activity, making it likely that we were recording from the same population of neurons as those described in other studies of area LIP (Barash et al., 1991a, b; Thier et al., 1991). Barash et al. (1991a) suggested that LIP neurons with tonic memory activity and pre-saccadic activity may be involved in the planning of saccades. We found that most neurons with spatially tuned auditory and visual memory activity also had pre-saccadic responses in both the auditory and the visual tasks, lending support to a saccade-planning role for these neurons.

Anatomical considerations

The connections of PPC with auditory areas have not been studied in detail. There are connections from area 22/area TA (areas AA1-3 of Pandya and Yeterian, 1985)—the auditory association cortex—to area 7 (both the surface and the inferior bank of the intraparietal sulcus; Divac et al., 1977). Area Tpt (the temporoparietal junction) receives connections from parakoniocortex (paAlt) (Pandya and Sanides, 1973), considered to be part of the auditory association cortex (Pandya and Yeterian, 1985), and projects on to area 7 (Pandya and Kuypers, 1969). Leinonen et al. recorded in this area auditory responses that seemed selective for sound source location. There are also connections from the superior temporal polysensory area of Bruce et al. (1981) to area 7a (Andersen et al., 1990a) and LIP (Baizer et al., 1991). Baylis et al. (1987) have reported many auditory responses in single units in the dorsal superior temporal sulcus (areas TS and TAa). Since the various subdivisions of the PPC are densely interconnected (Pandya and Seltzer, 1982; Andersen et al., 1990a) it seems reasonable to assume that auditory information can gain access to the whole PPC, but indirectly.

Functional role of area LIP

Auditory responses have been described before in the PPC. In early studies Hyvärinen et al. and Mountcastle et al. (1975) tested several PPC neurons with a few auditory stimuli, such as the jingling of keys and hand-clapping, and reported no responses. The neurons tested may have been outside area LIP, which had not yet been identified as a separate area. Various other authors reported auditory responses in portions of PPC (Koch and Fuster, 1989, Sakata et al., 1973; Seal et al., 1983). Interestingly, these authors only found responses to auditory stimuli when they were cues for movement. We cannot know whether the same is true for the responses we recorded in LIP because in the task we used the stimulus was always a cue for movement.

Previous recording and lesion studies of area LIP have focussed on its role in visually guided tasks. Goldberg et al. (1977, 1990, 1992; Duhamel et al., 1992a) have interpreted the activity of LIP neurons as important for visual sensory and attentional processing. Gnadt et al. (1988) and Barash et al. (1991a, b) systematically studied the responses of LIP neurons in monkeys performing visual memory saccades. The responses they observed led them to hypothesize a role for area LIP in sensorimotor integration to guide eye movements. In such a scheme the stimulus-related responses would encode a stimulus' location. The sensory activity would give rise, within the same cells and/or in other cells, to memory activity that would encode the metrics of the upcoming saccade. Saccade centers downstream of LIP could use this signal to generate the appropriate saccade.

Our findings establish that area LIP is not only concerned with processing visual stimuli. We believe we recorded from the same population of neurons as Barash et al. (1991a, b) because we observed response properties in the visual memory saccade task largely similar to the ones they reported (specifically, latencies of the S responses, co-occurrence of S and M responses, proportions of excitatory and inhibitory responses and of contralateral and ipsilateral spatial preferences). Thus area LIP contains a population of neurons that

responds to stimuli of at least two modalities. These results extend the possible roles of area LIP in the processing of sensory stimuli. Rather than being restricted to processing retinal events, LIP neurons appear to integrate sensory cues of multiple modalities to encode the spatial location of a relevant stimulus. This integration is consistent with the sensorimotor processing necessary to program orienting movements, regardless of the modality calling for such movements.

FIGURE LEGENDS

Figure 1:

Activity of a neuron with auditory and visual S-period responses in the auditory and visual memory saccade tasks. In each of the four panels time is plotted on the abscissa. The rows of ticks indicate the occurrences of spikes, one row per trial. Below these is a histogram indicating the average firing rate across different trials. The two traces below the histogram represent the horizontal (E_h) and vertical (E_V) components of eye position. The spike rasters, histogram and eye position traces for each trial are horizontally aligned on stimulus onset. The double arrows above panel A indicate the stimulus (*S*) and memory (*M*) periods of the memory saccade task. The thick horizontal lines below each panel indicate the presentations of the stimuli: *FP* = fixation point; *Sound R* = right speaker; *Sound L* = left speaker; *Light R* = light spot on the right; *Light L* = light spot presented on the left. A. Auditory memory saccade to the right. B. Auditory memory saccade to the left. C. Visual memory saccade to the right. D. Visual memory saccade to the left. Scales are 100 ms/hor. div., 10 (impulses/sec)/vert. div. (firing rate) and 15 deg./vert. div. (eye position).

Figure 2:

Activity of a neuron with auditory and visual S-period and M-period responses. All panels and their labels are as in Fig. 1. A. Auditory memory saccade to the right. B. Auditory memory saccade to the left. C. Visual memory saccade to the right. D. Visual memory saccade to the left.

Figure 3:

Proportions of the cells with auditory S- or M-period responses that are in various categories of response types. Each bar shows the number of cells in a response category as a percentage of the cells in our database.

The *dark bars* indicate major response categories, while the *light bars* indicate subcategories of the dark bars just above them (the total number of cells in a group of subcategories is the same as the number of cells in the category marked by the dark bar just above them).

Figure 4:

Latencies of onset of S-period responses of neurons with clear S responses. A. Latencies of auditory S responses. B. Latencies of visual S responses. C. Differences between auditory and visual latencies (auditory - visual) for cells with both auditory and visual clear S responses.

Figure 5:

Activity of a neuron with directionally tuned saccade-related activity during (A) a memory saccade cued by the left speaker, and (B) a memory saccade cued by a light on the left. The spike rasters, firing rate histogram and eye position are plotted as in Fig. 1, except that all events are horizontally aligned on the beginning of the saccade. Scales are as in Fig. 1.

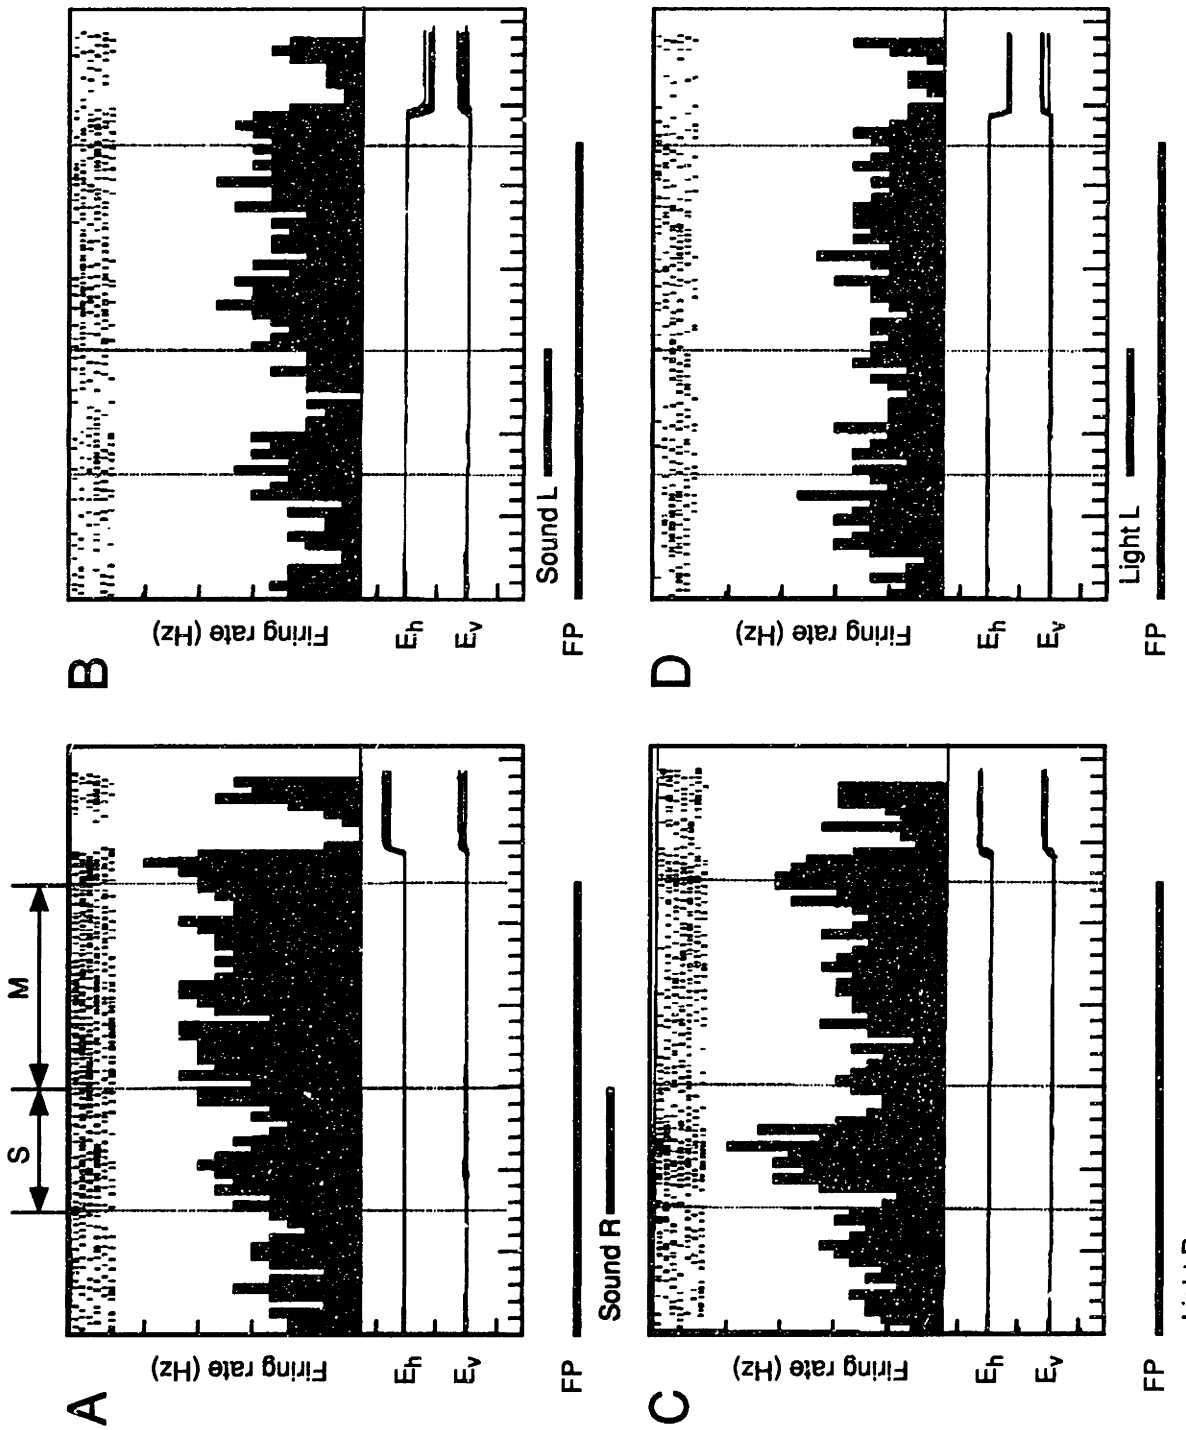

Figure 1

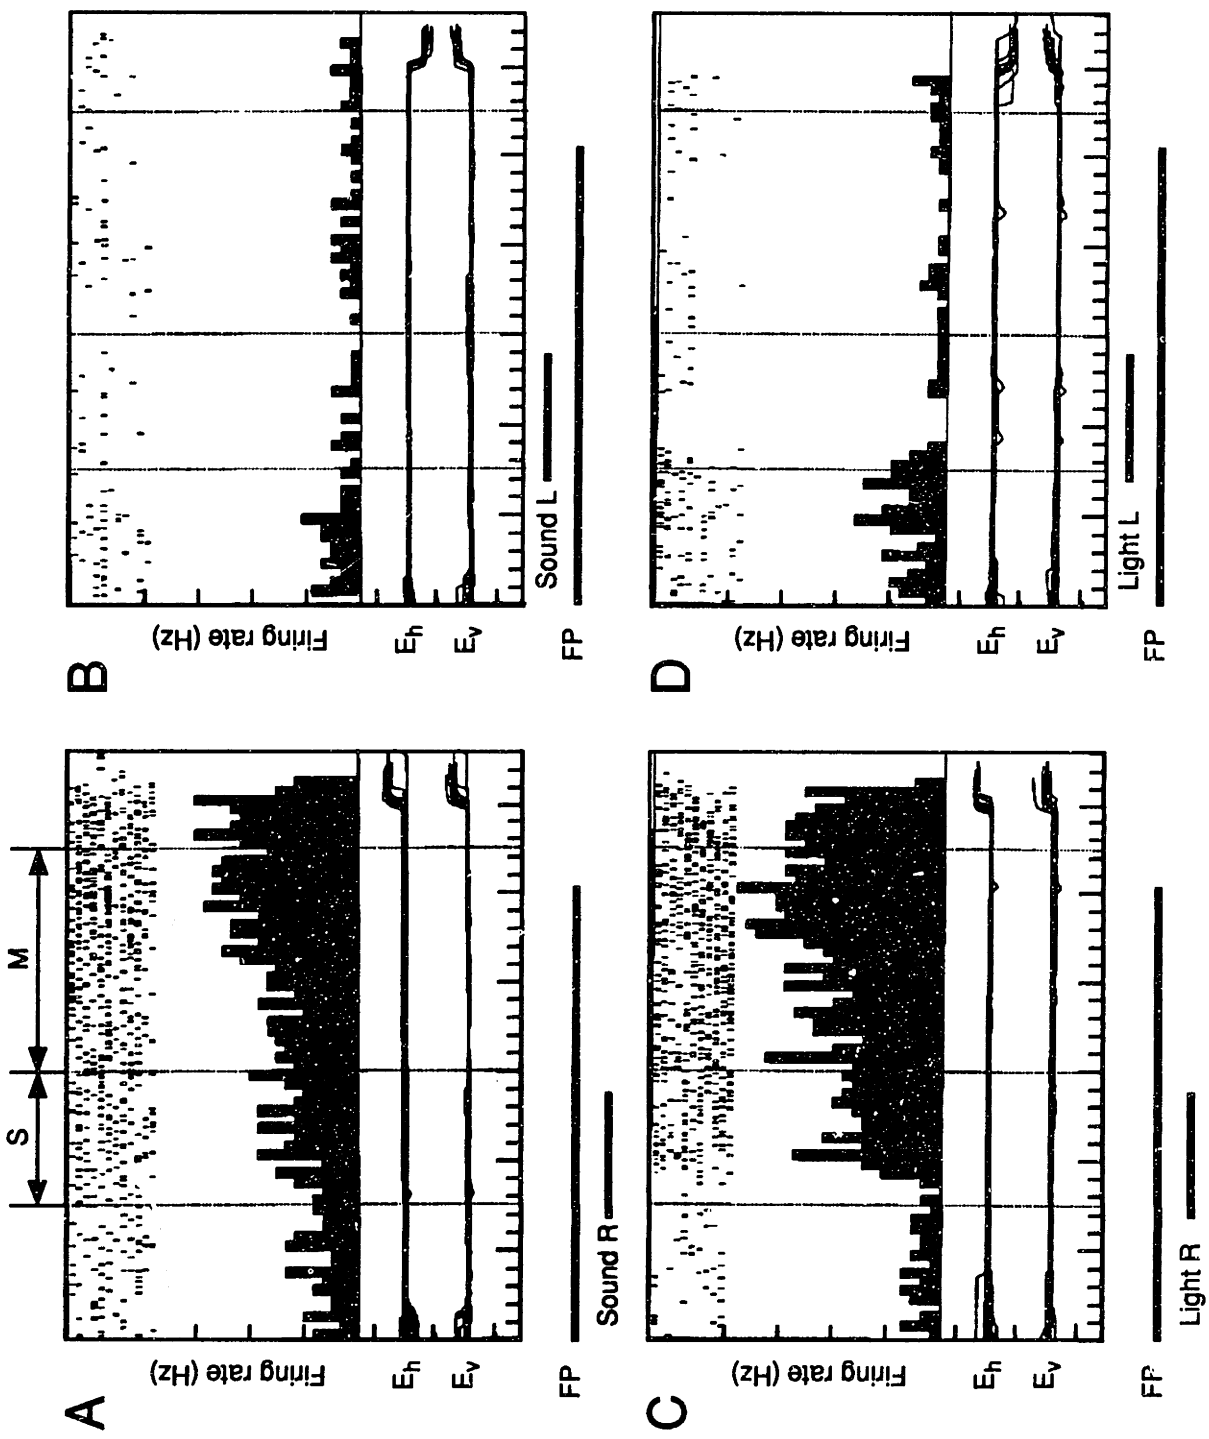

Figure 2

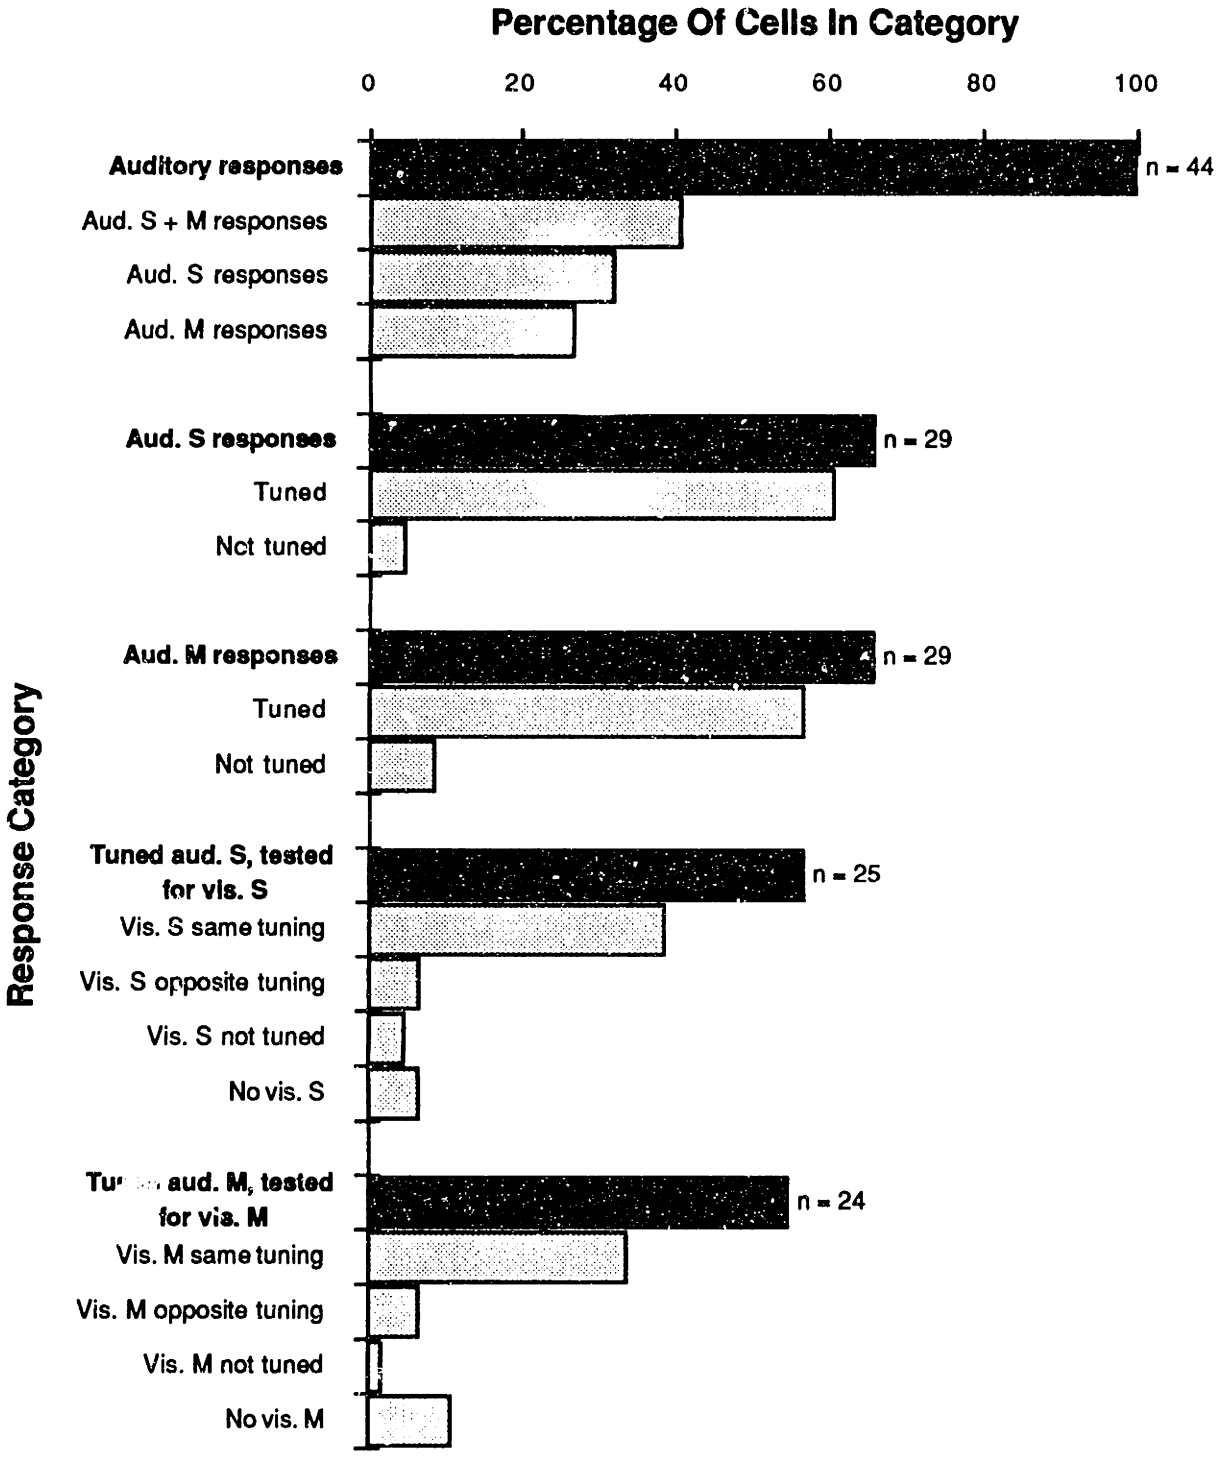
Percentage Of Cells In Category
0
20
40
60
80
100
Response Category
Auditory responses
n = 44
Aud. S + M responses
Aud. S responses
Aud. M responses
Aud. S responses
n = 29
Tuned
Not tuned
Aud. M responses
n = 29
Tuned
Not tuned
Tuned aud. S, tested for vis. S
n = 25
Vis. S same tuning
Vis. S opposite tuning
Vis. S not tuned
No vis. S
Tuned aud. M, tested for vis. M
n = 24
Vis. M same tuning
Vis. M opposite tuning
Vis. M not tuned
No vis. M

Figure 3

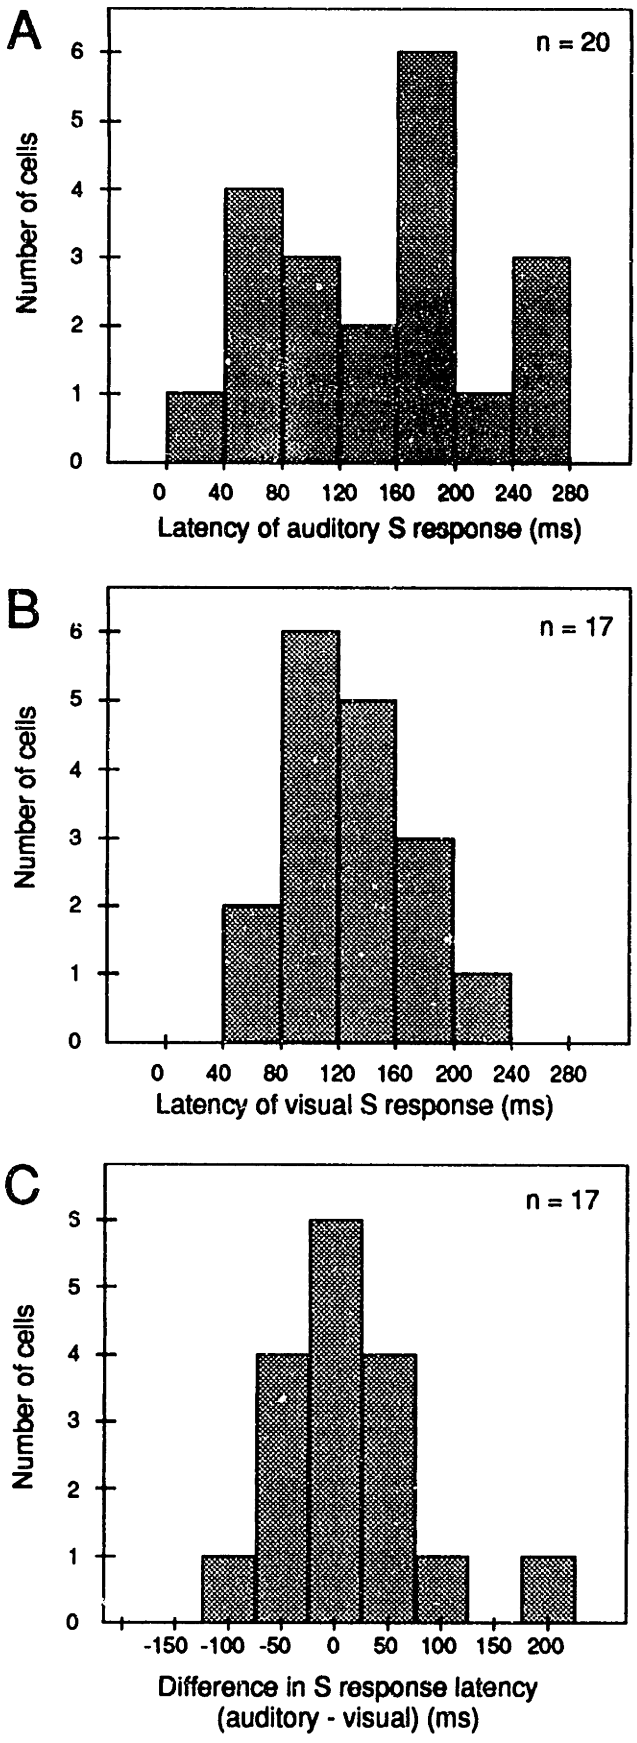
A
n = 20
Number of cells
0
1
2
3
4
5
6
0 40 80 120 160 200 240 280
Latency of auditory S response (ms)
B
n = 17
Number of cells
0
1
2
3
4
5
6
0 40 80 120 160 200 240 280
Latency of visual S response (ms)
C
n = 17
Number of cells
0
1
2
3
4
5
6
-150 -100 -50 0 50 100 150 200
Difference in S response latency
(auditory - visual) (ms)

Figure 4

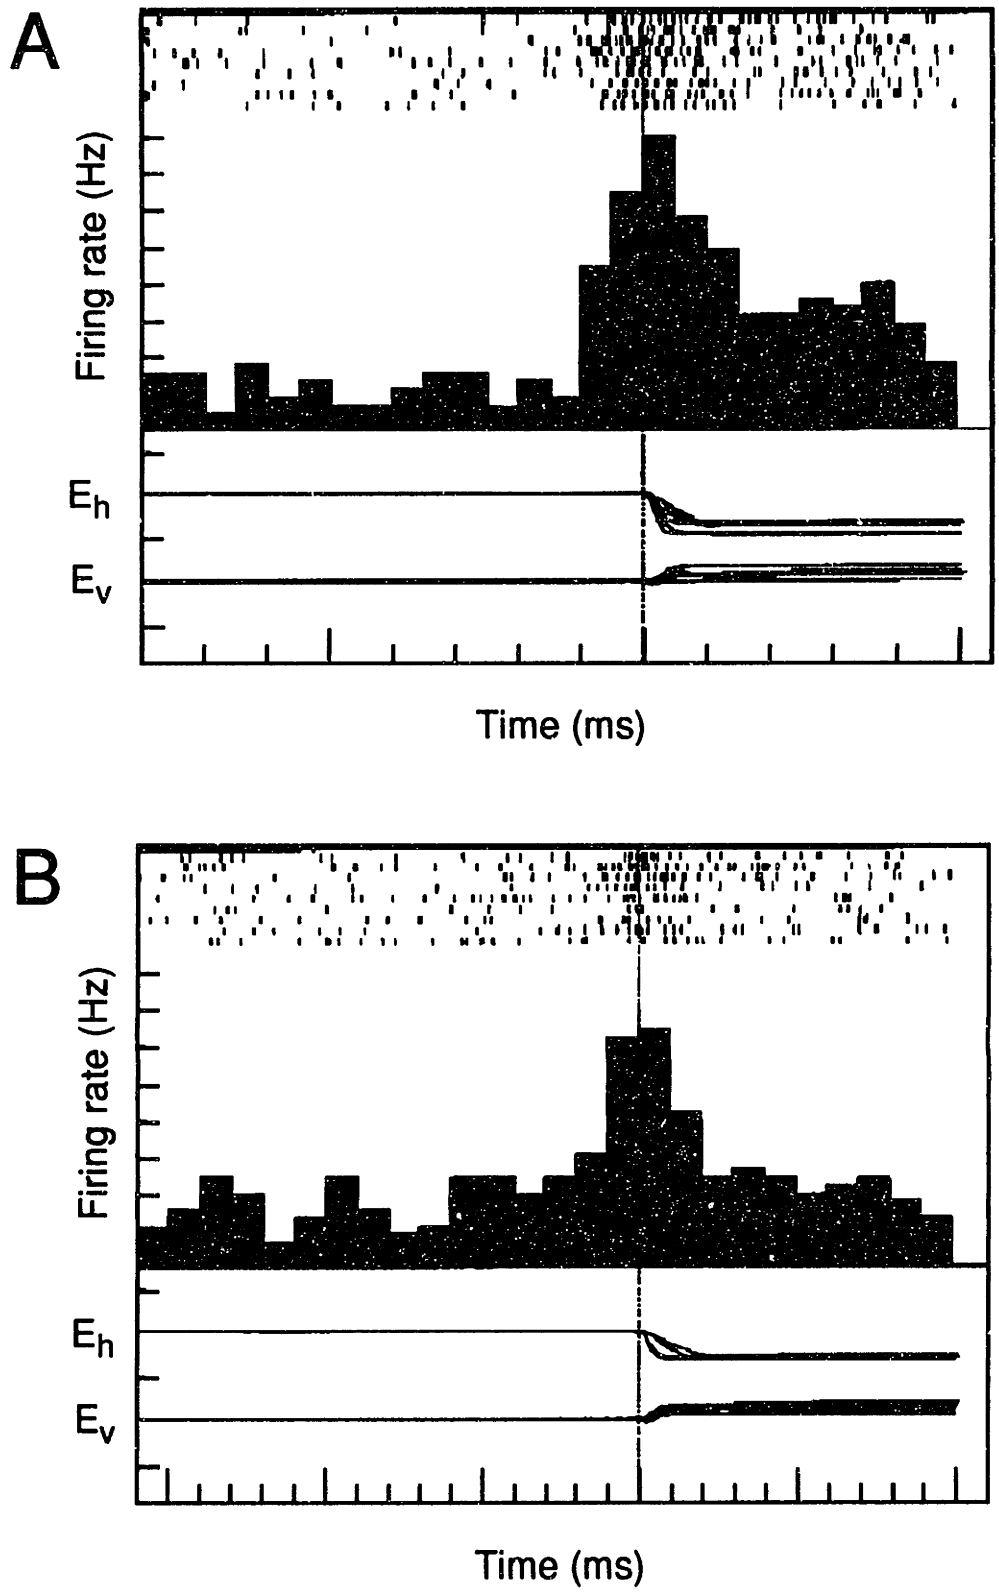

Figure 5

Chapter 6

Looking ahead: In what coordinates are sound locations encoded in area LIP?

SUMMARY

We saw in the previous chapter that LIP neurons respond as the monkey prepares saccades to acoustically cued locations. Primates localize auditory stimuli using cues, such as interaural intensity differences and interaural phase differences, that are anchored to a coordinate frame centered on the head. The auditory system of many vertebrates contains several representations of auditory space in head-centered coordinates. The fact that area LIP neurons have spatially tuned auditory responses raises the question of what coordinate frame is used by these neurons to encode the location of a sound source. If this area integrates spatial cues from various sensory modalities, then it must transform the visual and auditory cues into a common reference frame. This reference frame might be centered on the head, which would require no transformation of the auditory signals. It might be anchored to the eye, requiring a transformation of the auditory signals into oculomotor coordinates. Finally, the common reference frame might be an intermediate one between purely craniocentric and purely oculocentric.

We report in this chapter the initial results of experiments addressing this question. Our approach is to record the activity of LIP neurons while a monkey makes memory saccades to acoustic targets starting from various initial eye positions. Changing the initial eye position should have no effect on the neurons' responses if these are encoded in head-centered coordinates, while it should alter the responses if these are in eye-centered coordinates. We have found that the auditory responses of some LIP neurons are indeed affected by eye position. These effects are varied but suggest that auditory and visual spatial maps can be brought into register in area LIP.

INTRODUCTION

We have shown in the previous chapter that the lateral intraparietal area (area LIP) of the posterior parietal cortex (PPC) contains neurons that respond to auditory stimuli. These neurons are—at least coarsely—sensitive to the location of a sound source that is the target of an upcoming saccade. This spatial sensitivity may imply that these neurons encode the location of acoustic stimuli or the metrics of the saccade necessary to direct gaze at a sound source (see also Chapters 7 and 8). Both possibilities raise the issue of the coordinate frame in which these spatial responses are encoded.

The auditory system of primates extracts the location of a sound source by comparing differences in intensity and arrival time of sounds reaching the two ears (Rayleigh 1876; Blauert, 1983; Yost and Gourevitch, 1987; Syka and Masterton, 1987). These physical cues, being obtained via the ears, are anchored to a reference frame attached to the head. The simplest possibility is thus that the auditory signals of area LIP area encoded in a head-centered reference frame. It is also possible that these signals are transformed to a different coordinate frame. A transformation might be expected for the same reasons that visual signals may be expected to undergo coordinate transformations: because the head can move relative to the rest of the body, the location of a sound source may need to be transformed to another coordinate frame in order to program appropriate orienting or reaching movements toward the sound source (see also Chapter 2). The PPC is *a priori* a candidate structure for such transformations to occur for reasons analogous to why it was a candidate structure for the transformation of visual stimuli before recording studies addressing this question (Andersen and Mountcastle, 1983; Andersen et al. 1985b, 1987, 1990b; Barash et al., 1991a, b) were undertaken: lesions of the PPC have been observed to impair the ability to localize and orient to auditory stimuli in humans and monkeys (Heilman et al., 1970, 1971; Heilman and Valenstein, 1972; see also Chapter 4), though not as frequently or clearly as for the visual modality.

There is a more compelling reason for asking in what coordinate frame are auditory responses encoded in area LIP. As we saw in Chapter 5 most LIP neurons with auditory responses also have visual response. Of these bimodal neurons, moreover, most of the ones that are spatially tuned have matching spatial sensitivities (at least to the extent of left vs. right preferences), suggesting that they encode spatial locations and/or metrics of the planned saccade independently of stimulus modality. The coordinate frame in which these responses are encoded would thus be a feature essential to their role as a spatial signal. We would expect the coordinate frame of the visual and auditory responses of an individual neuron to be in register if these responses are to signal a unique spatial location. Otherwise as the eyes moved in the orbits a mismatch would arise between the neural representation of the location of a seen visual stimulus and the location from which the same stimulus emitted sounds. Therefore if the auditory and visual spatial representations established by LIP neurons were found not to be in register, one would have to suspect either a functional role of these neurons other than spatial coding, or the existence of other neural signals that could disambiguate area LIP's auditory spatial code from its visual one.

The question of spatial coordinates of auditory signals and their relation to those of visual signals has been explored before—for reasons similar to the ones outlined above—in the superior colliculus (SC) of the monkey (Jay and Sparks, 1984, 1987a, b). Neurons in the deep layers of this structure respond to various cues to the location of a saccade target and encode various parameter of the metrics of saccades (for a review see Sparks, 1989; but see also the recent results of Stanford et al., 1994). Many deep SC neurons can be activated by visual, auditory, or somatosensory stimuli. These neurons encode the vector of the saccade necessary to foveate the stimulus independently of the stimulus' modality. Jay and Sparks (1984, 1987a, b) thus examined whether the auditory and visual responses of multimodal neurons in the monkey's SC were in spatial register. They did so by recording neural responses to a sound stimulus as the eye's direction of

fixation was varied. The auditory responses of most neurons was affected by the initial eye position. These responses were thus *not* in head-centered coordinates. This interaction, moreover, was such as to shift the neuron's auditory receptive field (RF) toward an eye-centered topography. The auditory RF's of most neurons, as recorded from different initial fixation positions, aligned to one another much more closely when plotted against oculomotor error (corresponding to the vector of the saccade required to foveate the stimulus from the given fixation position) than when plotted against the stimulus' position in head-centered space.

We have begun a series of experiments to examine the coordinate frames of auditory responses in area LIP and their relationship to the coordinates of visual responses. The first step in these experiments is to establish whether auditory responses are in the same coordinate frame as the physical cues to sound location—i.e. a head-centered frame—or are at all transformed. Any effect of varying the eyes' fixation position on the responses would imply a transformation away from a head-centered code. We are using methods similar to those of Jay and Sparks (1987a) to answer this question for LIP neurons. We have trained a monkey to make memory saccades to auditory stimuli (described in Chapter 5) from different initial fixation positions. We have found neurons with auditory responses that are affected by the initial angle of gaze as well as neurons that are not affected. The neurons with an effect of eye position appear to transform the coordinates of the auditory stimulus into a variety of frames. At this point the small number of neurons does not permit us to draw any general conclusion regarding a single or general transformation of auditory signals occurring in area LIP. Some of the cells we have recorded so far suggest interesting processes that the auditory signals seem to undergo in this area. We will describe examples of what types of effect we have observed so far and discuss the types of transformations that may underlie these effects.

METHODS

The methods of this study are the same as those described in Chapter 5 except as noted below. Briefly, we trained a rhesus monkey (*Macaca mulatta*) to make memory saccades to auditory and visual stimuli in total darkness. We then prepared the monkey for chronic recording of single unit activity and recorded the activity of neurons in area LIP of the PPC while the monkey made memory saccades to auditory stimuli from various fixation positions. For details see Chapter 5.

The difference between the experimental paradigm of this study and the one in Chapter 5 lies in the arrangement of stimuli guiding the animal's behavior. As shown schematically in Fig. 1a, the fixation spot could appear at one of three locations (marked in Fig. 1 by a cross) separated horizontally by 12 deg. of visual field angle. The center spot was positioned straight ahead at approximately 10 deg. above eye level. Five speakers were arranged below the three possible fixation spot locations, approximately at the monkey's eye level. These were spaced horizontally at 12 deg. intervals, the center speaker being in the monkey's median plane.

The timing of the memory saccade task is shown in Fig. 1b. At the beginning of a trial the fixation spot was turned on directly at one of its three possible locations (left, center, and right), and the monkey started fixating on it. After 750 ms of fixation one of the speakers emitted a noise burst for 500 ms. The monkey had to continue fixating for another 500 ms after stimulus offset (memory period, M). At this point the fixation spot was extinguished and the monkey had to make a saccade, in the dark, to the remembered location where the stimulus had appeared.

Trials starting from any of the three fixation positions and using any of the five speakers were pseudorandomly interleaved—that is, the fixation position-speaker combination was chosen randomly among all possible ones for a given trial, but each combination that was performed successfully was removed from the choice of possible

combinations until all combinations were exhausted. This method ensured the sampling of an equal number of trials for each fixation point-speaker combination.

The monkey did not perform memory saccades with a horizontal component larger than 24 deg. reliably enough to allow consistent data collection in an interleaved-trial paradigm. Therefore we collected data for saccades to: all speakers from the central fixation point; all but the rightmost speaker from the left fixation point; and all but the leftmost speaker from the right fixation point. The monkey's performance varied from day to day but was above 75% correct on average.

We computed the average firing rate during all but the first 100 ms of the M period. The first 100 ms were excluded to avoid possible contamination by the tail end of responses related to the stimulus.

No histology has yet been performed on the monkey used in this experiment. Assignation of recorded neurons to area LIP thus remains tentative and is based on site and depth of electrode penetration and the presence of clear directionally tuned saccade-related activity.

RESULTS

Predictions

We examined the effect of initial eye position on the M-period activity of the auditory memory saccade task by plotting this activity against two variables. In each plot grouping the responses according to the initial fixation position. One variable is the target's position, that is, the horizontal position of the speaker on the screen. Because the monkey's head was fixed, the target's position corresponded to the speaker's head-centered coordinates. If a neuron has auditory responses encoded in head-centered coordinates, then a plot its activity vs. target position should consist of three overlapping curves. Each curve describes the neuron's response to sounds coming from different directions (the neuron's auditory RF). Changing the eye position should not affect head-centered coordinates, so the three curves

representing the neuron's RF recorded at different eye positions should be identical. If there is any effect of eye position, on the other hand, then the three curves would look different. Any deviation of the three curves from one another would imply an effect of eye position on the auditory responses, and thus a deviation of the neuron's spatial responses from purely head-centered coordinates.

The other variable against which we plotted the responses is motor error. This is the vector of the saccade required to foveate the speaker's location from the given initial eye position. These coordinates are the speaker's eye-centered coordinates, corresponding to the retinal position, relative to the fovea, of the speaker. In the case that the coordinates of auditory responses are transformed, one can ask whether they are transformed to an eye-centered reference frame. This is the reference frame in which responses to visual stimuli are encoded in many areas of the brain. Plotting the auditory responses of LIP neurons against oculomotor error thus allows us to compare the coding of these responses to previously established coding of visual responses in this and other brain structures.

Effect of eye position on auditory responses

Some of the neurons we found encoded speaker locations in head-centered coordinates. Fig. 2 shows the response patterns observed for one such neuron. This cell has spatially tuned auditory M responses, which area strongest for sounds from the center speaker (at 0 deg. in Fig. 2a). As the eye position is varied (solid vs. dashed vs. dotted curves in Fig. 2a) the neuron's response to each speaker does not vary significantly. In this plot the responses for two or three eye positions are plotted at each target position. These roughly overlap at each point, the peak being always at the 0 deg. speaker location. The neuron thus has a RF in purely head-centered coordinates. The auditory signal has not undergone any spatial transformation.

In Fig. 2b the same three response curves are plotted as in Fig. 2a, but this time versus oculomotor error. The same RF functions as in Fig. 2b appear shifted horizontally by 12 deg. relative to one another.

This is exactly what we expect given that the curves are aligned in the target position plot. The lack of alignment means that the RF's are definitely not in eye-centered coordinates.

Other neurons showed a systematic effect of eye position. The neuron whose responses are plotted in Fig. 3 has an auditory RF with a single peak, just as the neuron in Fig. 2 did. When recorded with the eye starting from the central fixation point this RF showed a peak at the 12 deg. speaker location (Fig. 3a, solid line). When eye position was varied, however, the neuron's gave different responses to the same set of speakers (Fig. 3a, dashed and dotted lines). It appears from Fig. 3a that the neuron's RF maintains the same shape but is shifted mostly horizontally. In particular, the RF seems to shift by about 12 deg. to the left when the eye is moved 12 deg. to the left, and it shifts by about 12 deg. to the right when the eye moves by 12 deg. in the same direction. Such shifts suggest a reference frame attached to the eye, which is confirmed by the plot of activity vs. motor error (Fig. 3b). In this plot the three curves show very strong alignment, with similar responses from any eye position for a given saccade amplitude and a peak for 12 deg. rightward saccades. The coordinates of this neuron's auditory signals are systematically affected by eye position in such a way as to produce eye-centered coding of the speaker's location.

Another interesting type of interaction of eye position and auditory responses is shown in Fig. 4. Like the neurons described above, this neuron has a spatially graded auditory RF with a single peak (Fig. 4a, solid line). As for the neuron of Fig. 3, the peak of this neuron's RF, when plotted in head-centered coordinates, shifts as the eye position changes (Fig. 4a). The shift is of the same direction and magnitude as the eye position shift, which suggests that the RF is in eye-centered coordinates. The horizontal shift of the RF peak, however, is not the only change occurring as the eyes move. The overall response level also changes, increasing as the fixation point moves to the right. The plot of activity vs. motor error (Fig. 4b) thus reveals an eye-centered RF with a peak at 12 deg. rightward saccades that is modulated in amplitude by changes in eye position.

We also found other types of interaction between auditory responses and eye position (not shown). Notable among these were head-centered RF's modulated in amplitude by eye position, and RF's that shifted with the eye when the eye moved into one hemifield but not the other.

DISCUSSION

The initial results of this study have revealed a clear effect of eye position on the M-period responses of some LIP neurons during the auditory memory saccade task. In some neurons the interaction is systematic and effects a transformation of the auditory responses into oculomotor coordinates. In other neurons eye position modulates the responses' amplitudes, while either leaving the RF in head-centered coordinates or also shifting it to eye-centered coordinates. In yet other neurons the interaction of eye position and auditory responses does not produce a code that can be easily summarized or interpreted.

Retinal vs. oculomotor coordinates

In the previous chapter we have shown that many area LIP neurons show spatially tuned stimulus-period and memory-period responses during an auditory memory saccade task. Responses in the same task periods of visually cued memory saccades have been extensively studied (Gnadt et al., 1988; Andersen et al., 1990b; Barash et al., 1991a, b). The stimulus period responses in the visual version of the task contain information regarding the location of the visual stimulus and can thus be described as sensory responses. Previous studies (Gnadt et al., 1988; Barash et al., 1991a, b) and those described in Chapters 7 and 8, on the other hand, have established that the M-period responses of most LIP neurons encode the vector of the next planned saccade, while those of a small proportion of neurons encode a memory trace of the stimulus' location. We report in Chapter 7, moreover, that even the stimulus-period responses of some LIP

neurons contain information related to the planned saccade, and thus cannot be considered solely sensory responses.

Such distinctions between the sensory and motor planning roles of the auditory stimulus- and memory-period responses have not been studied. In the study of Chapter 5 the sensory and motor components of a memory saccade were separated temporally but not spatially. The vector of the stimulus' location relative to the fovea (retinal coordinates) always matched the vector of the saccade required to foveate the stimulus (oculomotor coordinates). Until these issues are addressed experimentally we must keep this ambiguity in mind as we speculate on the possible roles of the eye position effects we have observed in this study.

For the same reasons we refer to RF's that align on the motor error plot as "eye-centered." They could specifically be in retinal or oculomotor coordinates, but we cannot distinguish between these possibilities in the paradigm we have used.

Transformation of the coordinates of visual stimuli in area LIP

A previous study established that the responses of most LIP neurons elicited in the visual memory saccade task are affected by eye position (Andersen et al., 1990b). For most of these neurons the response fields remained aligned in an eye-centered reference frame. The modulation by eye position mostly affected the amplitude of the response, producing response gain fields that made the encoding of stimulus location (or saccade vector) ambiguous in the activity of single neurons. Modelling studies had shown, however, that the gain field code could underlie a representation of a stimulus' head-centered coordinates distributed over the activity of a population of neurons (see Chapter 2 for a review). Area LIP can thus transform the coordinates of a visual stimulus from a retinal (eye-centered) reference frame to a head-centered one. More recent studies have shown that this transformation does not stop at the head-centered representation. The activity of a group of LIP neurons is also modulated by head position in a gain field fashion, producing a distributed representation of stimuli

in a trunk-centered reference frame (Brotchie and Andersen, 1991), while some LIP neurons appear to encode stimulus locations in external (world-centered) coordinates (Snyder et al., 1993).

All this evidence implies an important role of area LIP in the transformation of the coordinates of visual stimuli. What is clear is that there is not a unique final reference frame into which all LIP neurons transform stimulus coordinates. The response patterns observed using visual stimuli can subserve multiple representations of locations in various coordinate frames. These representation may then be used separately or in concert for the purpose of programming various types orienting movements—such as saccades, head and eye gaze shifts, and trunk rotations—, allocating spatial attention, reaching, etc. It is in light of these multiple possible functions of area LIP that we would like to briefly speculate on the possible roles of the response patterns we observed in this study.

Possible roles of the observed transformations

The auditory M responses of neurons unaffected by eye position (such as in Fig. 2) may serve to simply encode the locations of auditory stimuli in a head-centered reference frame, without any transformation from their original encoding. This signal is analogous to the purely retinotopic code of visual stimulus location observed in a minority of LIP neurons (Andersen et al., 1990b), in the sense that it is in the same reference frame as the physical cues reaching the sensory organ. This signal could be matched to the distributed code of the head-centered location of visual stimuli—for example, in order to correctly localize a sound and an image from the same source. It could alternatively serve as an input signal to a network of neurons combining this signal with eye position information in order to transform it into eye-centered coordinates (to produce the encoding shown in Figs. 3 and 4). Finally, the head-centered signal could serve as input to a network of neurons that combined it with head position signals to produce a representation in body-centered coordinates.

The neurons with M responses in purely eye-centered coordinates (such as in Fig. 3) carry an auditory signal completely transformed into an unambiguous code of locations relative to the eye. No analogous transformation (in the reverse direction) of visual signals—that is, a single-neuron code of visual stimuli in head-centered coordinates—has been observed in area LIP. This signal may work in conjunction with the purely eye-centered code of visual stimuli of a minority of LIP neurons, possibly to encode the next planned saccade in oculomotor coordinates instead of head-centered ones. Indeed, such auditory and visual signals may coexist in single neurons—such as the ones described in Chapter 5—and would encode saccade vectors unambiguously at the single-neuron level.

The neurons with M responses modulated in amplitude by eye position (such as in Fig. 4) seem to have gain fields analogous to those of the majority of LIP neurons tested in visual memory saccades (Andersen et al., 1990b). The ones that also have receptive fields in eye-centered coordinates (such as in Fig. 4) encode the location of an auditory stimulus (or the goal of the next planned saccade) in a manner that may be indistinguishable from—or at least equivalent to—the representation of visual stimuli observed for the majority of LIP neurons (Andersen et al., 1990b). As described above, this code consists of eye-centered RF's modulated in amplitude by eye position, and produces a distributed code of head-centered location. Such transformations of visual and auditory cues would allow individual neurons to encode locations independently of stimulus modality. It may be the way most of area LIP solve the problem posed by the mismatch between the original coordinate frames of visual and auditory stimuli: neither by completely transforming the eye-centered coordinates of images into head-centered ones, nor by completely transforming the head-centered coordinates of sounds into eye-centered ones; but by introducing an eye position signal that produces a partial transformation for each modality. The location code in such a representation becomes ambiguous at the single-cell level but is

resolved into a head-centered representation at the neuronal population level.

An obvious next step in these experiments is to record the eye position effects on both the visual and the auditory responses of LIP neurons in order compare the transformations to which visual and auditory stimuli are subjected in the same neurons. Such studies would elucidate whether any of the schemes hypothesized above is actually implemented in area LIP, as well as clarify how some of the problems posed by the integration of two different sensory modalities may be solved in this part of the brain.

FIGURE LEGENDS

Figure 1:

Experimental paradigm. *a*) Spatial layout of the stimuli. The crosses indicate the three positions where the fixation light spot could appear. Their horizontal spacing is 12 deg. Below the fixation points is a row of five speakers, also separated horizontally by 12 deg. *b*) Outline of the events of a memory saccade. First the fixation point (FP) appears and the monkey starts to maintain gaze on it, as shown in the idealized eye position trace (eye pos.). After 750 ms the stimulus (in this experiment one of the speakers) is turned on for 500 ms (sensory period). The FP remains on for another 500 ms, after which it is turned off and the monkey makes a saccade, in darkness, to the remembered location of the stimulus.

Figure 2:

Tuning curves of an LIP neuron encoding sound location in head-centered coordinates. The plots show the M-period activity (mean ± s.e.m.) for auditory memory saccades to each speaker from each of the three fixation points. Each line connects the set of responses obtained while the monkey maintained gaze on one fixation position. Solid line: central fixation; dashed line: left fixation; dotted line: right fixation. *a*) The responses are plotted against the position of the speaker on the screen (head-centered coordinates). *b*) The responses are plotted against horizontal motor error, that is, the horizontal component of the saccade vector required to look at the speaker from the given fixation point.

Figure 3:

Tuning curves of an LIP neuron encoding sound location in eye-centered coordinates without amplitude modulation. Plot conventions are the same as in Fig. 2.

Figure 4:

Tuning curves of an LIP neuron encoding sound location in eye-centered coordinates with modulation of response amplitude by eye position. Plot conventions are the same as in Fig. 2.

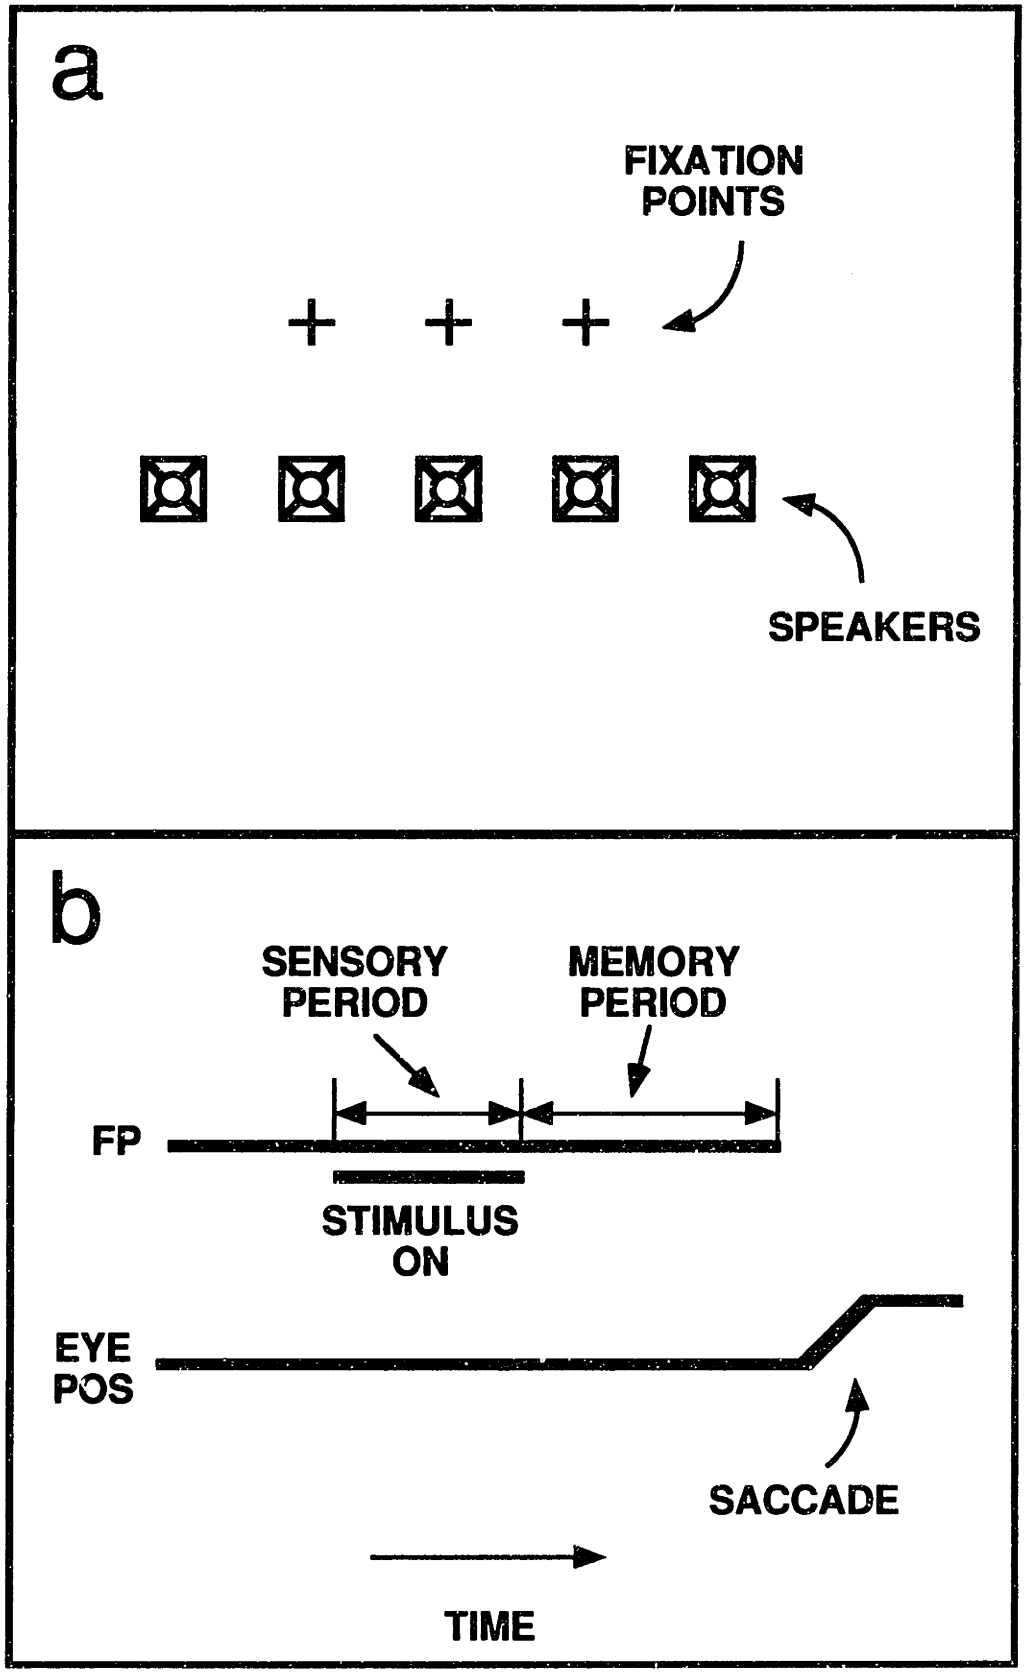
a
FIXATION POINTS
SPEAKERS
b
SENSORY PERIOD
MEMORY PERIOD
FP
STIMULUS ON
EYE POS
SACCADE
TIME

Figure 1

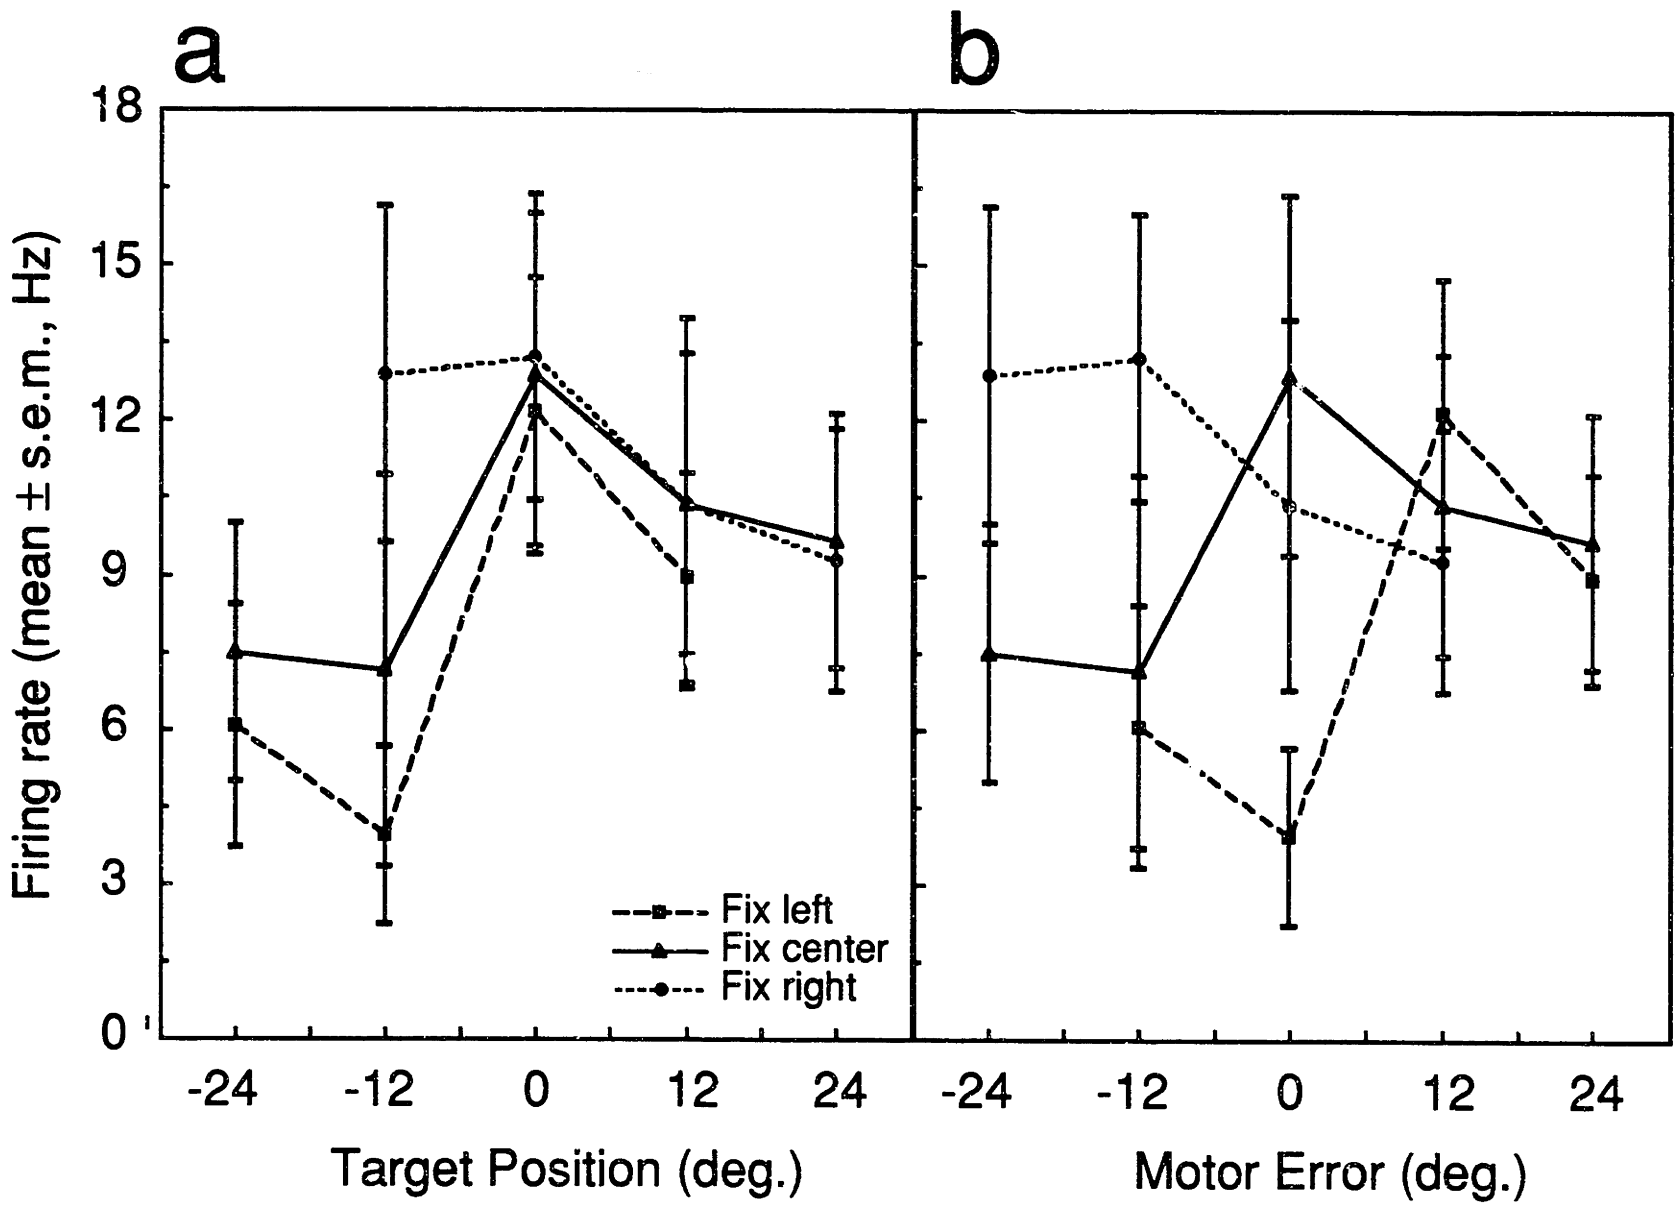
a
b
18
15
12
9
6
3
0
Firing rate (mean ± s.e.m., Hz)
Fix left
Fix center
Fix right
-24
-12
0
12
24
Target Position (deg.)
Motor Error (deg.)

Figure 2

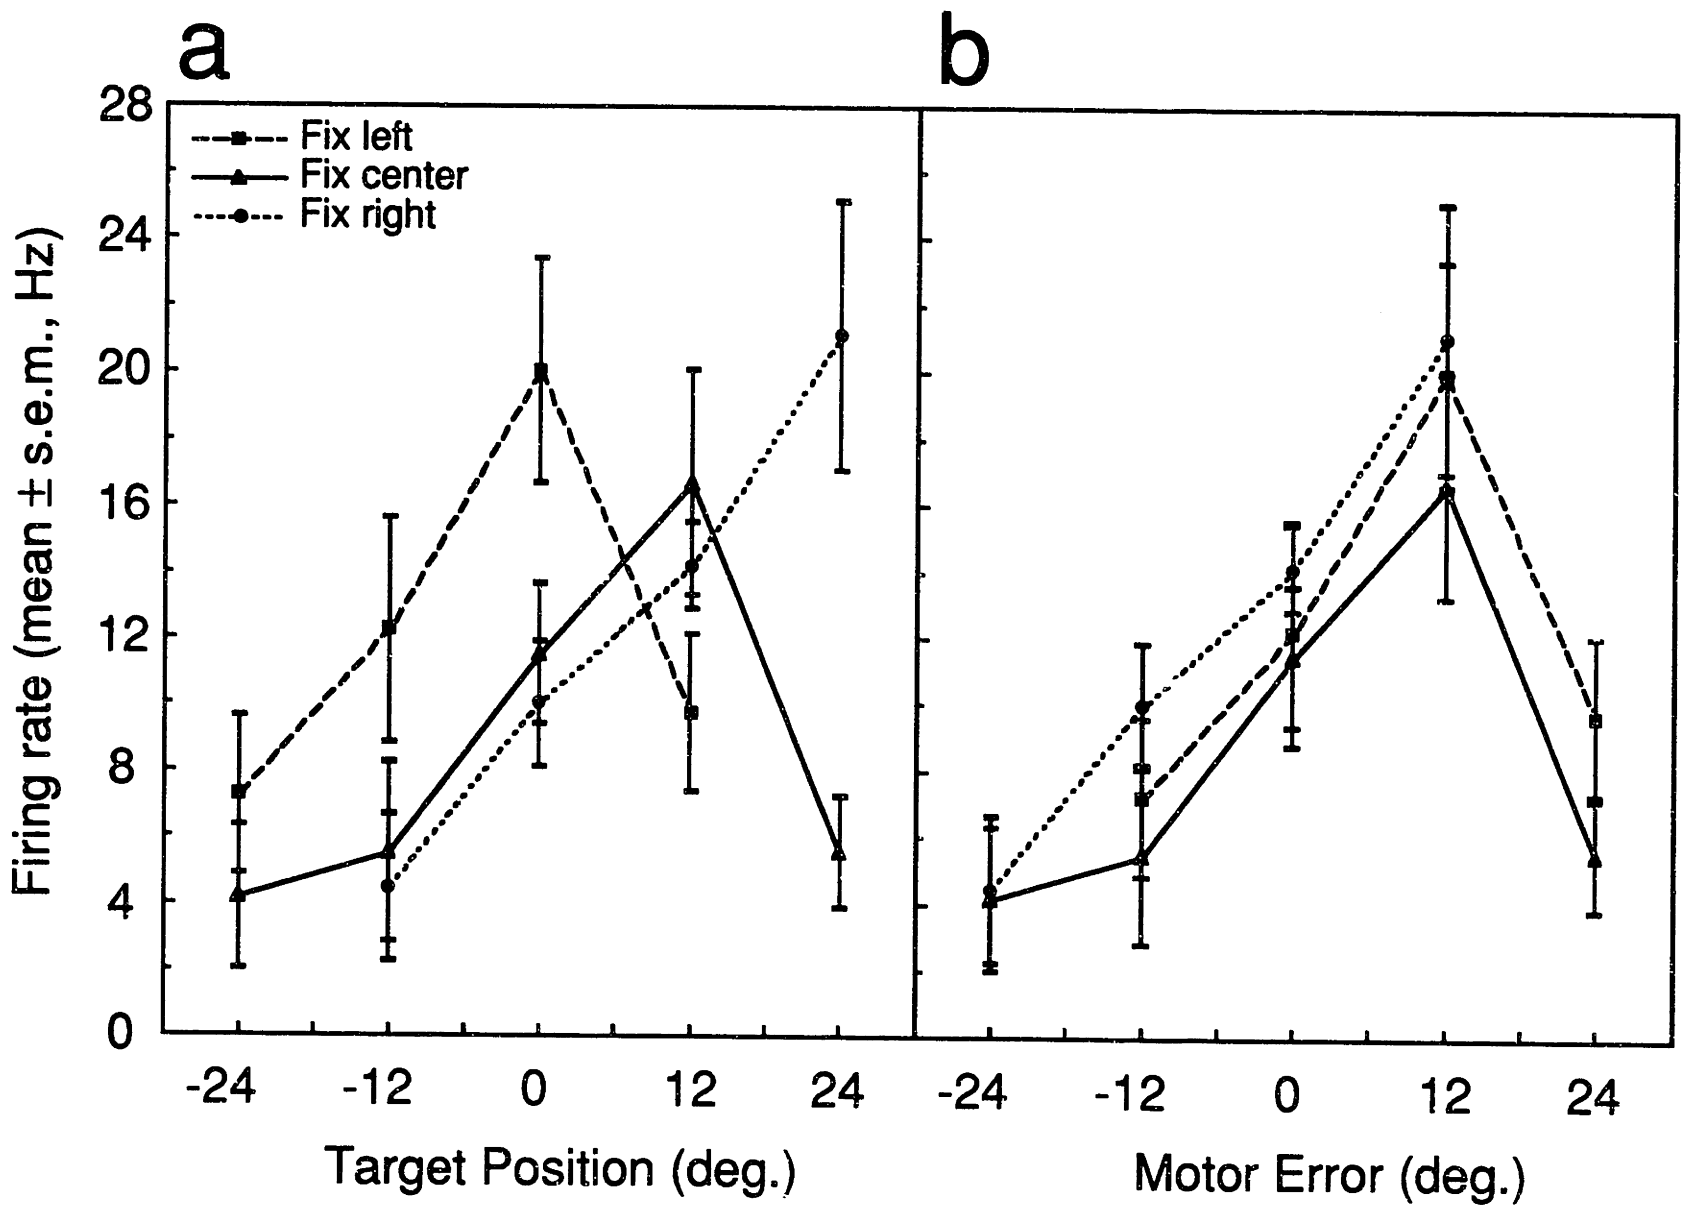
a
b
Fix left
Fix center
Fix right
Firing rate (mean ± s.e.m., Hz)
0
4
8
12
16
20
24
28
-24
-12
0
12
24
Target Position (deg.)
Motor Error (deg.)

Figure 3

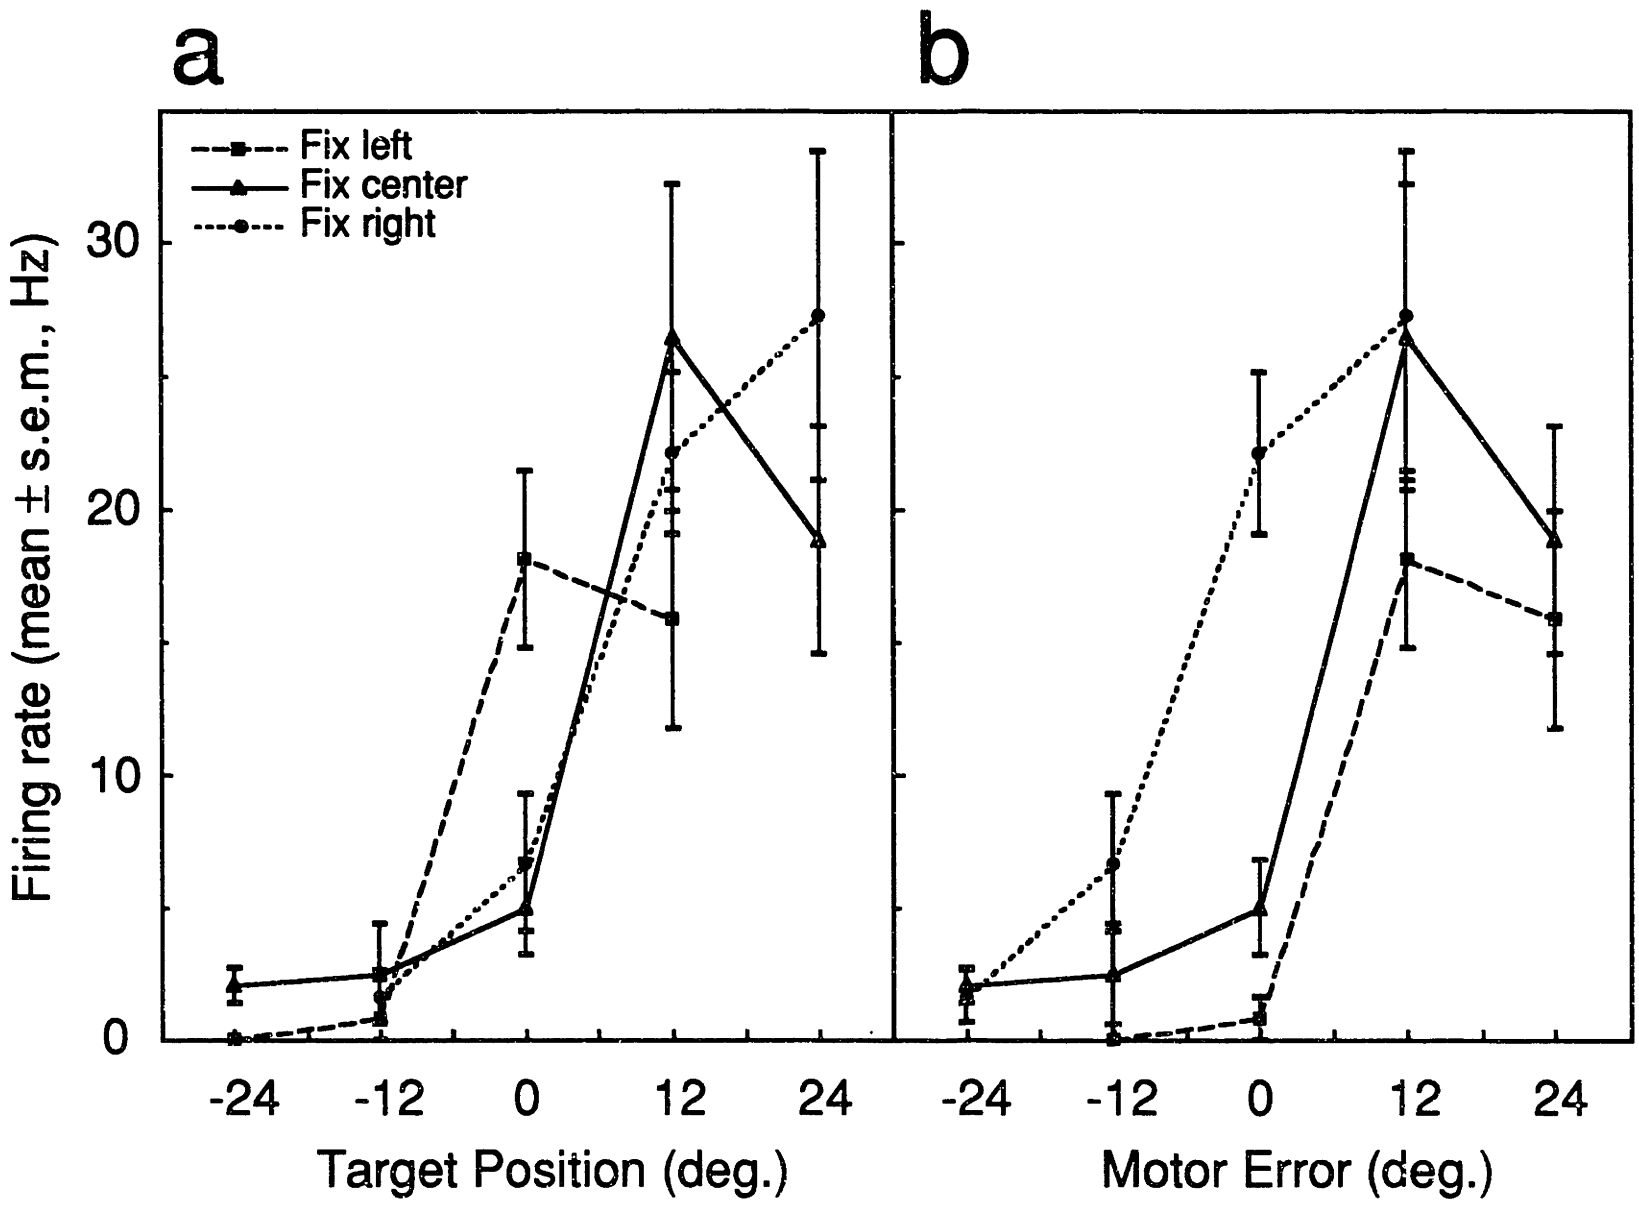
a
b
Fix left
Fix center
Fix right
Firing rate (mean ± s.e.m., Hz)
0
10
20
30
-24
-12
0
12
24
Target Position (deg.)
Motor Error (deg.)

Figure 4

III

FORMING A MOTOR PLAN

Chapter 7

Motor intention activity in area LIP. I. Dissociation of motor plan from sensory memory

SUMMARY

The lateral intraparietal area (area LIP) of the monkey's posterior parietal cortex (PPC) contains neurons that are active during saccadic eye movements. The activity of these neurons includes a component related to the visual stimulus and one related to the saccade (Mountcastle et al., 1975; Andersen et al., 1987). The visual and saccade responses are spatially tuned and the location of a neuron's visual receptive field (RF) relative to the fovea generally overlaps its preferred saccade amplitude and direction (PD) (Barash et al., 1991a). When a delay is imposed between the presentation of a visual stimulus and a saccade made to its location (memory saccade task), many LIP neurons maintain elevated activity during the delay (memory activity, M) (Gnadt and Andersen, 1988; Andersen et al., 1990; Barash et al., 1991a; reviewed in Andersen et al., 1992). Previous studies suggest that the M activity codes the metrics of the next intended saccadic eye movement; that is, the activity is a memory trace of what the animal intends to do (Gnadt and Andersen, 1988; Barash et al., 1991b). Recent studies have alternatively suggested that LIP neurons encode the locations of visual stimuli (Goldberg et al., 1990; Duhamel et al., 1992) regardless of where the animal intends to look. In this study we examine whether the memory activity of LIP neurons specifically encodes movement intention (*motor plan* hypothesis), or whether it also encodes the locations of recent visual stimuli (*sensory memory* hypothesis). In the companion study (described in the next paper, Bracewell et al., 1994) we examine whether the intended-movement activity reflects changes in motor plan; that is, when the animal selects a different target for a movement, whether the movement vector encoded in area LIP shifts to reflect this change in plan.

We trained monkeys (*Macaca mulatta*) to memorize the locations of two visual stimuli and plan a sequence of two saccades, one to each remembered target, as we recorded the activity of single LIP neurons. Two targets were flashed briefly while the monkey maintained fixation; after a delay the fixation point was extinguished and the

monkey made two saccades in sequence to each target's remembered location, in the order in which the targets were presented. This "delayed double saccade" (DDS) paradigm allowed us to dissociate the location of visual stimulation from the direction of the planned saccade and thus distinguish neuronal activity related to the target's location from activity related to the saccade plan. By imposing a delay we eliminated the confounding effect of any phasic responses coincident with the appearance of the stimulus and with the saccade.

We arranged the two visual stimuli so that in one set of conditions at least the first one was in the neuron's visual receptive field (RF), and thus the first saccade was in the neuron's preferred direction (PD). M activity should be high in these conditions according to both the sensory memory and motor plan hypotheses. In another set of conditions the second stimulus appeared in the RF but the first one was presented outside the RF, instructing the monkey to plan the first saccade away from the neuron's PD. If the M activity encodes the motor plan it should be low in these conditions, reflecting the plan for the first saccade (away from the PD). If it is a sensory trace of the stimulus' location it should be high, reflecting stimulation of the RF by the second target.

We tested 54 LIP neurons (in 3 hemispheres of two monkeys) with M activity on the DDS task. Of these, 44 (81%) had M activity that specifically encoded the next intended saccade. They were active in the delay period, as expected, if the first saccade was in their preferred direction. They were less active or silent if the next saccade was not in their preferred direction, even when the second stimulus appeared in their RF.

The M activity of 8 (15%) of the remaining neurons was consistent with the sensory memory hypothesis. Their firing rate during the delay reflected stimulation of the RF independently of the saccade being planned. All but two of these cells did not show any activity between the first and second saccade when the second saccade was towards their RF. Most of the cells with sensory memory activity thus encode the location of the second stimulus but do not remap the

location of a sensory stimulus after an eye movement. The remaining two neurons had M activity that was not consistent with either hypothesis.

We also recorded the activity of a subset of neurons (n=42) in a condition in which no stimulus appeared in a neuron's RF but the second saccade was in the neuron's PD. In this case the majority of neurons tested (25/42, 60%) became active in the period between the first and second saccade, even if neither stimulus had appeared in their RF. Moreover, this activity appeared only after the first saccade had started in all but one of these neurons. The neurons' response fields thus were not predictively remapped in advance of saccades.

The majority of LIP neurons encode the next intended saccade. Their activity is not solely determined by visual stimulation. It does not reflect attention to a spatial location, either, because it was reduced or absent when a location required attention but not a saccade to it. Area LIP appears to be the first station along the occipito-parietal cortical visual pathway to express the intention to execute a movement.

INTRODUCTION

The posterior parietal lobe of the primate brain has been implicated in a variety of functions subserving sensorimotor integration. Certain regions of the posterior parietal cortex (PPC) seem especially important for the production of saccadic eye movements. Lesions of these regions in humans and monkeys impair the perception of spatial relationships in the visual field and the ability to make voluntary saccades (e.g., Balint, 1909; Hécaen and De Ajuriaguerra, 1954; Holmes, 1918; Lynch, 1980; Lynch and McLaren, 1989). Neurophysiological studies in awake behaving monkeys have revealed single unit activity in the PPC related to saccadic eye movements (Hyvärinen and Poranen, 1974; Mountcastle et al., 1975; Lynch et al., 1977).

Initially there was a controversy as to whether the activity occurring around the time of a saccade was a motor command (Mountcastle et al., 1975) or rather an artifact of sensory stimulation (Robinson et al., 1978). Later studies (Andersen et al., 1987) addressed this issue by recording the activity of posterior parietal neurons in a "delayed" or "memory" saccade task (introduced by Hikosaka and Wurtz, 1983, in studies of the basal ganglia). In this task a peripheral visual stimulus appears briefly while a monkey maintains fixation on a light spot; after a delay the monkey the fixation spot is turned off, which instructs the monkey to make a saccade, in the dark, to the location where the stimulus appeared. The memory saccade paradigm separates temporally the sensory and motor components of the saccade task. The initial studies showed that PPC neurons often carry both visual and saccade-related signals (Andersen et al., 1987). The exact role these signals played in the production of saccades, however, remained unclear.

The visual and saccade-related signals are especially prominent in the lateral intraparietal area (area LIP), a subdivision of the PPC characterized by strong projections to eye movement centers (especially the frontal eye fields, FEF, and the superior colliculus, SC) as well as

multiple inputs from other extrastriate visual areas (Lynch et al., 1985; Andersen et al., 1990a; Blatt et al., 1990). The responses of neurons in this area have been characterized using the memory saccade task (Andersen et al., 1990b; Barash et al., 1991a, b). These signals are spatially tuned. Visual responses vary across the visual field, being strongest for stimuli in a circumscribed sensory response field (receptive field, RF). Saccade-related responses are broadly tuned for amplitude and vary with saccade direction, reaching a maximum for saccades in the neuron's preferred direction (PD). The spatial tuning of the visual and saccade-related responses in LIP generally coincide, that is, the RF is in the same direction, relative to the fovea, as a saccade in the neuron's PD.

Besides responding during the visual stimulus' presentation and during the saccade, many LIP neurons maintain sustained activity during the delay period of a memory saccade (Gnadt et al., 1988; Andersen et al., 1990b; Barash et al., 1991a). This "memory" (M) activity has similar spatial tuning to the visual and oculomotor responses (Barash et al., 1991b). It could reflect a memory of the stimulus' location, a covert shift of attention within the visual field, or the intention to execute the upcoming saccade.

The studies by Gnadt and Andersen (1988) and Barash et al. (1991b) showed that the responses of LIP neurons are coded in oculomotor coordinates. Using a double saccade paradigm these experiments showed that LIP activity appears before a saccade made in the neuron's preferred direction even without RF stimulation. These authors thus proposed that the M activity is a memory trace of what the animal intends to do.

Other studies of area LIP have offered another interpretation of the role of this area in sensorimotor integration (Goldberg et al., 1990; Duhamel et al., 1992a). According to these studies the major role of area LIP is to construct a perceptual map of visual space by encoding the locations of visual stimuli and maintaining this representation anchored to a retinally based reference frame across eye movements (Goldberg et al., 1990; Duhamel et al., 1992a). Neural activity in this

area would thus indicate that a stimulus is or has been at a particular location in the visual field, independently of whether the animal wants to look at that location.

In this study we examined whether LIP neurons specifically encode movement intention (*motor plan* hypothesis), or whether their memory activity also encodes the locations of visual stimuli (*sensory memory* hypothesis). To answer this question we extended the double saccade task used by Gnadt and Andersen (1988) and Barash et al. (1991a, b). In a double saccade task (first described by Hallett and Lightstone, 1976, and applied to experiments in monkeys by Mays and Sparks, 1980) two peripheral targets are presented in very fast sequence while the monkey fixates. The monkey must then make a sequence of two saccades to the locations of the two targets. By choosing appropriate locations of the targets relative to a neuron's RF one can tease apart the relationships of the neural activity to the locations of sensory stimuli and to saccade metrics.

Because in a simple double saccade task the monkey makes the eye movements as soon as possible, the visual and saccade-related responses cannot be separated (the reaction time before the first saccade being on the order of 150 ms). We thus added a delay requirement to this task. In the "delayed double saccade" (DDS) task, two visual stimuli appeared in sequence at different locations while the monkey maintained fixation, and were followed by a delay (Fig. 1). After the delay the fixation point was extinguished and the monkey had to make two saccades, in darkness, to the remembered location of each stimulus. During the delay period the monkey had to remember the locations of two visual stimuli and plan a saccade to the location of the first stimulus and then to that of the second one. Because in the delay period he was maintaining fixation, we could observe neural activity underlying sensorimotor integration uncontaminated by sensory or motor events.

By varying the locations of the stimuli relative to a neuron's RF the location of sensory stimulation can be dissociated from the metrics of the saccade being planned. Gnadt and Andersen (1988) and Barash et

al. (1991b) arranged the two stimuli so that they both fell outside the neuron's RF, but so that the second saccade was in the neuron's PD. Most LIP neurons became active between the first and second saccade, showing that sensory stimulation is not required to elicit LIP responses and that these responses predict the upcoming saccade vector.

In order for LIP activity to *specifically* encode the plan for the next saccade (as in the motor plan hypothesis), it should 1) appear every time the monkey prepares to make a saccade in the neuron's PD, whether the stimuli were inside or outside the RF (as Gnadt and Andersen, 1988 and Barash et al., 1991b showed); and 2) be reduced or absent when the next saccade is away from the PD, even if the RF is stimulated. Alternatively LIP responses could encode the presence of a stimulus or the allocation of attention to a salient location in the visual field (sensory memory hypothesis), as Goldberg et al. (1990) and Duhamel et al. (1992a) have suggested. In this case, these responses should appear every time a salient target (such as a saccade target) appears in the neuron's RF, regardless of whether the monkey plans the next saccade to that location or not.

We employed five different arrangements of stimulus locations (shown in Fig. 2, which will be described in more detail below). In two of these arrangements (classes 1 and 3) the first visual target fell in the RF and the first saccade was in the PD. These stimulus classes established a neuron's M response when the saccade planned was towards the neuron's RF (*congruent conditions*).

In two other classes (classes 4 and 5) the first target was outside the RF and the second one was inside the RF. The first saccade was thus away from the PD in spite of stimulation of the RF (*incongruent conditions*). The motor plan hypothesis predicts that M activity should be absent or reduced in the incongruent condition relative to the congruent condition. The sensory memory hypothesis, on the other hand, predicts that M activity should be similar in both conditions.

Finally, in one class both stimuli were presented outside the RF but were aligned so that the second saccade was in the neuron's PD. The

motor plan hypothesis predicts that the neuron will become active in the period between the first and second saccade, reflecting the plan for the next saccade in its PD, while the memory hypothesis predicts that the neuron will remain silent throughout the trial.

In the experiments described in this chapter the M activity of most LIP neurons reflected the next planned saccade and not the location of the sensory stimulus. In the next chapter we show that this activity appears when the monkey plans to make a saccade in a certain direction independently of whether the saccade is actually made. Area LIP activity therefore reflects a monkey's intention to make a specific saccade. These results suggest that the parietal lobe plays a role not only in the analysis of the sensory world but also in the preparation for movement.

METHODS

Animals, surgery, and animal care

We used two adult male rhesus monkeys (*Macaca mulatta*) in this study. We prepared each monkey for chronic recording of eye position and cortical neural activity through three surgical procedures. These were conducted with the monkey under general anesthesia (10 mg/kg intramuscular ketamine followed by 10 mg/kg intravenous pentobarbital sodium (Nembutal) using aseptic techniques. In the first procedure we implanted a scleral search coil in one eye (Judge et al., 1980; Robinson, 1963) and mounted a metal head post in dental acrylic on the skull. In two separate procedures we implanted a recording chamber on each hemisphere over the posterior parietal cortex (Brodmann's areas 5 and 7). After each procedure the monkeys received analgesics and systemic antibiotics and rested for a week.

We trained the monkeys via operant-reinforcement techniques in several saccade tasks including the ones used in this study. During the training and recording periods the monkeys were deprived of water in their home cages and received apple juice or water as reward for correct task execution, supplemented by additional water at the end of

each session to reach the required daily ration. They had at least two days of rest per week with unrestricted access to water. The monkeys received routine veterinarian care, and their well-being was observed in accordance with National Institutes of Health guidelines.

Experimental set-up and data collection

The monkey sat in a completely dark room facing a large featureless tangent screen placed 57 cm away. Small light spots (c. 0.5 deg. diameter, c. 45 cd/m^2) were back-projected onto the screen from two projectors through galvanometer-controlled mirrors. A laboratory computer (Digital Equipment Corp. PDP 11/73) presented the stimuli and monitored the monkey's behavior. We sampled eye position at 500 Hz using the scleral search coil method (Judge et al., 1980; Robinson, 1963) and we recorded extracellularly the action potentials of single cortical neurons with glass-coated Pt-Ir microelectrodes (Wolshbart et al., 1960) mounted on a Chubbuck microdrive. The computer stored the eye position samples and the time of occurrence of action potentials for off-line analysis.

Behavioral tasks

Each monkey learned to perform several tasks involving saccades for the purposes of several studies. The ones used in this study are the memory saccade task and the delayed double saccade task.

A *memory saccade (MS)* trial started when a spot was turned on directly in front of the monkey, at eye level, and the monkey started fixating on it. After 800 ms of fixation a peripheral stimulus was presented for 300 ms. The monkey had to continue fixating for another 400 ms after stimulus offset (M period). At this point the fixation spot was extinguished and the monkey had to make a saccade, in the dark, to the remembered location where the stimulus had appeared. The stimulus was placed at an eccentricity of 5-25 degrees along one of 8 directions (the 4 cardinal and the 4 diagonal directions).

The responses of most LIP neurons during a MS trial consist of at least one of three components. These are a sensory response (LS,

appearing during light stimulus presentation), a saccade-related response (SR, coincident with saccade execution), and sustained activity during the delay between stimulus presentation and fixation spot offset (memory activity, M) (Gnadt et al., 1988; Barash et al., 1991a, b). These signals are spatially tuned. LS responses vary across the visual field, being strongest for stimuli in a circumscribed sensory response field (receptive field, RF). SR responses vary with saccade direction (and to some extent with amplitude), reaching a maximum for saccades in the neuron's preferred direction (PD). The spatial tuning of the visual and saccade-related responses in LIP generally coincide, that is, the RF is in the same direction, relative to the fovea, as a saccade in the neuron's PD. The spatial tuning of the M activity generally matches that of the LS and SR responses (Barash et al., 1991b). LIP neurons thus have up to three spatially tuned response fields (sensory, memory, and motor), which are generally aligned. We used MS trials to locate each neuron's visual RF and saccadic PD.

If a neuron had sustained M activity in MS trials with targets at 10 or 20 deg. eccentricity we then tested it in the *delayed double saccade (DDS)* paradigm. The DDS task consisted of up to 5 classes of trials, all involving a fixation point placed directly in front of the monkey at eye level and two peripheral stimuli at different locations. The timing of the stimuli was the same in all classes (Fig. 1a). A trial started when the monkey began fixating on the FP straight ahead. After a period of simple fixation (400 or 500 ms) the first visual stimulus (target 1, T1) appeared and was followed, after a brief interstimulus interval, by the second stimulus (T2). T2 was followed by a delay (the M1 period) during which the animal continued to maintain fixation. At the end of this delay the FP was extinguished and the monkey had to make two saccades (S1 and S2) in darkness, first to the remembered location of T1 and then to the remembered location of T2. For most cells the targets were presented after 500 ms of fixation, for 50 ms each, separated by 50 ms, and followed by an M1 period of 500 ms. For a few cells the targets appeared after 400 ms of fixation, for 200 ms each, separated by 200 ms, and followed by an M1 period of 400 ms.

The monkeys also made pauses of variable lengths between the first and second saccade. We refer to the period between the two saccades as M2.

We used five classes of DDS stimuli, each having a particular arrangement of the saccade targets relative to the neuron's RF (Fig. 2). In class 1 both targets (T1, T2) fall in the neuron's RF and the first saccade (S1) is in the neuron's PD. In class 2 neither target falls in the RF and the second saccade (S2) is in the PD. In class 3 only the first target falls in the RF and the first saccade is in the PD. In class 4 only the second target stimulates the RF and the second saccade is in the PD. Note that in classes 3 and 4 the visual stimuli are at the same locations but are shown in opposite sequence. In class 5 the second target is in the RF but neither saccade is in the PD. The possible patterns of neuronal activity for each of these classes will be described below.

We tested neurons in class 5 only if their M-response tuning curve (as estimated from M saccade trials) was narrow enough so that the first stimulus of DDS class 5 was truly outside the neuron's M response field.

Data analysis

We focused our analysis on neural activity during the delay periods of each task. For each neuron we computed the average firing rate during the delay period (except for the first 100 ms) and subtracted from it the average background firing rate (computed from 300 to 800 ms from the start of each trial for cells tested with an 800 ms fixation period, and from 100 to 400 ms for cells tested with a 400 ms fixation period). We defined the resulting net firing rate as a neuron's M response for each behavioral class and used two-tailed t tests to determine whether it was positive (excitatory), negative (inhibitory), or absent in each class. We also used t tests to compare the M responses of classes 1 and 5, and of classes 3 and 4 (see Results). For certain trial classes of several neurons we manually adjusted the boundaries of the M1 and background periods in order to better represent the neurons'

responses. If a cell's response clearly started well within one of these periods we redefined the period so that it consisted mostly of the cell's response time.

We also computed the net response of most neurons in the intersaccadic (M2) period of classes 2 and 4. We did this by aligning each trial on the beginning of the first saccade and then manually choosing a time segment that did not overlap with any part of the first or second saccade in any trial within that class. We took the difference between the average firing rate in this segment and the background firing rate as the M2 response. For several neurons the time between the two saccades was too short to compute an M2 response.

In order to describe the overall activity patterns of population of neurons in the various trial classes we computed for each neuron an activity index, I_a, based on its M1 response. This index was introduced by Barash et al. (1991a) and is defined as I_a = (average M1 firing rate - average background firing rate) / (average background firing rate)$^{1/2}$. For inhibitory cells we computed the absolute value of this index so that we could display index values of excitatory and inhibitory cells in the same plot. This index estimates the signal-to-noise ratio of a neuron's response. This index produces a measure of a neuron's response that is much less biased by its background firing rate than the difference in average firing rates alone, thus allowing comparison of responses across neurons. The M1 response is divided by the square root of the background firing rate, rather than by the background firing rate itself, to avoid assigning disproportionately small responses to neurons with inherently high background activity (see Barash et al., 1991a).

The responses and index values of a neuron in certain DDS classes were excluded from the analysis according to the following criteria. Because classes 1 and 2 tested the pattern of memory-period responses for stimuli located at 20 deg. eccentricity, these classes were excluded from the analysis for neurons that showed no M responses for stimuli at 20 deg. eccentricity (that is, no response in classes 1 *or* 2). Classes 3 and 4, on the other hand, tested the pattern of memory-period

responses for stimuli located at 10 deg. eccentricity. They were thus excluded from the analysis for neurons that had no M responses for stimuli at 10 deg. eccentricity (that is, no response in classes 3 *or* 4).

The latency of onset of the M2 response in class 2 was obtained by first computing the time histogram of a neuron's class 2 activity using 20 ms bins, aligning each trial with the beginning of the first saccade (defined as the time when the eye's tangential velocity became higher than 50 deg./s for at least 25 ms; see Barash et al., 1991a). We then compared via a *t* test the average firing rate in each 20 ms bin to the average firing rate during a baseline period within the delay period. This was from 300 to 0 ms before the saccade for most neurons, but was adjusted manually for a few neurons to include only a period during which the neuron was completely silent. The latency was defined as the lower bound of the first of two bins in which the activity was significantly different ($p<0.05$) from the activity in the baseline period. In order to avoid overestimating the latency of a few neurons whose M2 response developed gradually, we set the latency of these neurons at the first bin that was different from background at a p value <0.1, as long as this bin was followed by two bins that were significantly different at $p<0.05$.

Histology

The neurons described in this study were isolated in area LIP of the right and left hemispheres of two monkeys. After several months of recording the monkeys were euthanized and the neurons' locations were reconstructed based on each penetrations' chamber coordinates and depth relative to various landmarks. These landmarks consisted of several DC electrolytic lesions (made in the last few weeks before the monkey was euthanized), fluorescent dye injections, and guide wires inserted in the brain before sectioning it. The procedures we followed were the same as described in Barash et al. (1991a).

RESULTS

Database

Our database consists of 54 neurons isolated in area LIP in 3 hemispheres of two monkeys while the animals were performing the DDS task. These neurons are a subset of a large number of PPC neurons which we isolated for this experiment and for others that we performed in parallel (Barash et al., 1991a, b; Bracewell et al., 1991, 1994; Mazzoni et al., 1993). The neurons were selected according to the following criteria: first, assignation to area LIP (as described in the Methods section); second, presence of clear, spatially tuned M period activity (significant in class 1 of the DDS task, t test, $p<0.05$).

We have previously reported (Gnadt et al., 1988; Andersen et al., 1990b; Barash et al., 1991a, b) that the delayed saccade task allows us to distinguish three basic phases of activity in LIP cells: visual, delay period, and saccade-related. Fig. 3a illustrates the activity of a typical LIP cell while the monkey makes a memory saccade. There is a visual response (LS) that begins after the onset of the stimulus in the receptive field, then prolonged, sustained activity (M) during the delay period (during which there is no stimulus in the receptive field, and the monkey is not making any eye movements), and finally a second peak of activity (SR) occurring at the time of the saccade. Since the saccade is made is darkness to the remembered location of the target, the saccade-related response cannot be an artifact of visual stimulation. These findings have been described in detail by Barash et al. (1991a, b). Not all LIP neurons show all three phases of activity. In this study we further investigated the responses of units that exhibited clear M activity.

Fig. 3 illustrates another key aspect of LIP neurons' responses: they are spatially tuned. In Fig. 3a, target A is presented, 15 deg. above the fixation spot, and the neuron clearly shows LS, M and SR activity. However, when target B is presented, 15 deg. below the fixation spot, the cell has negligible activity in all phases of the trial (Fig. 3b). The LS, M and SR fields of any given neuron are typically

broad (c. 90 deg. width at half-maximal activity) but aligned with one another (Barash et al., 1991b). To be selected for further study in the present experiments, units had to show spatially selective M responses (the vast majority of LIP neurons were sufficiently narrowly tuned to meet this criterion).

All 54 neurons had spatially selective M activity (38 excitatory and 16 inhibitory) and were tested in DDS classes 1-4. To test a neuron in class 5 we further required that its M field be narrow enough so that the first saccade would clearly be outside its RF. Our sample for this class consists of 14 neurons.

Predictions

We designed the DDS task to distinguish between two hypotheses. During the first memory period (M1) the monkey must remember and attend to the locations of two sensory stimuli. Some amount of attention is presumably assigned to these locations as they are goals of future saccades. By attention we mean an enhanced allocation of perceptual resources to a selected locus in the visual field. Neural activity during the delay could reflect such processing of the locations cued by the stimuli (*sensory memory /attention* hypothesis). The monkey, on the other hand, is also planning the next saccade during this period. Neural activity could reflect some aspect of the formulation of this motor plan (*motor plan* hypothesis).

Because in the DDS task the metrics of sensory stimulation and planned saccade do not always coincide, we expect different response patterns based on whether or not the neural activity reflects motor planning processes. The response patterns predicted for the 5 DDS classes by each hypothesis are summarized in Fig. 2. If a neuron's M activity reflects a memory of the stimulus' location or a shift of attention to that location, then this activity should appear every time a saccade target appears in the neuron's RF, regardless of whether the monkey plans the *next* saccade to that location. In every trial the monkey must attend to and memorize the location of the second stimulus as well as the first one's. Therefore we should see activity in

the delay period (M1) of classes 1, 3, 4, and 5. If the M activity encodes, on the other hand, planning for the upcoming saccade, it should appear whenever the monkey prepares to make a saccade in the neuron's PD, whether the stimulus was inside or outside the RF. Thus we should observe activity during the M1 period of classes 1 and 3, and between the first and second saccades (period M2) in classes 2 and 4. M1 activity should *not* appear in classes 4 or 5, where the RF is stimulated but the first saccade (class 4) or both saccades (class 5) are not in the neuron's PD.

Activity in classes 1-4

The response pattern of an area LIP neuron in the first four classes of the DDS task is shown in Fig. 4. In class 1 both targets are in the neuron's RF and the first saccade is in its PD. The neuron responds to T1 with a high-frequency burst of spikes and then maintains sustained M activity until the first saccade is made (Fig. 4a). This activity is predicted by both the sensory memory and motor plan hypotheses (Fig. 2), and simply confirms the neuron's preference for stimuli in and saccades towards the lower left quadrant.

In class 2 neither stimulus falls in the RF. We see no activity in the M1 period, as predicted by both hypotheses. The neuron does become active, however, in the period between S1 and S2 (period M2; Fig. 4b). At this time the monkey has completed the first saccade and is preparing to make the second one. Activity in this period thus supports the motor plan hypothesis: the neuron becomes active before a saccade in its PD even in the absence of RF stimulation.

Classes 3 and 4 have identical spatial arrangements of stimuli. Because these appear in opposite order in the two classes, however, the saccade plans are different. In class 3 T1 is in the RF and S1 is in the PD. We see activity in M1 and not in M2, as predicted by both hypotheses (Fig. 4c). In class 4 T2 is in the RF and S2 is in the PD. Activity is absent in M1 but prominent in M2 (Fig. 4d). This response pattern supports the motor plan hypothesis because during M1 the planned saccade is opposite the PD, while S2 is in the PD. According

to the memory hypothesis the activity should have started in M1, immediately after stimulation of the RF by T2, and be maintained throughout M2, as the monkey had to remember the location of T2 throughout both M periods.

The M activity of the neuron of Fig. 4, when present, is of different magnitude in the various DDS classes. This can be explained by the fact that there are two sizes of retinal stimulation vectors and saccade vectors, one of 10 deg. and one of 20 deg. LIP neurons are often tuned for the amplitude, as well as the direction, of the stimulation and saccade vectors (Barash et al., 1991b). When we tested this neuron on trials of single memory saccades of different sizes, it produced higher M activity in the 20 deg. memory saccades than in the 10 deg. memory saccades (not shown). This tuning is reflected in the different amplitudes of responses in the DDS trials (Fig. 4). High M activity always precedes 20 deg. amplitude saccades (M1 in class 1, M2 in class 4), while lower (but still significant) M activity precedes 10 deg. saccades (M1 in class 2, M2 in class 3).

Activity in class 5

In class 5, as in class 3, the second target was in a neuron's RF. In contrast to class 3, however, T1 was positioned so that neither saccade would be in the PD. Class 5 thus complemented class 2 in contrasting the memory and plan hypotheses. In class 2 the RF was never stimulated but a saccade in the PD (S2) was planned. In class 5, conversely, the RF was stimulated but neither saccade was in the neuron's PD.

Fig. 5 shows the activity of an LIP neuron in this class. In trials of single memory saccades directed into the lower left quadrant the neurons showed clear M activity (Fig. 5a). In DDS class 5 trials (Fig. 5b) T2 elicits a response during its appearance. This response is suppressed, however, near the beginning of the delay period and remains suppressed throughout the execution of both saccades. There is no significant M activity during the memory period of class 5 (Fig. 5b), while it is clearly present in the memory period of the single

memory saccade (Fig. 5a). Thus as the monkey formulates his plan, during the delay period of class 5, to make the first saccade, the neuron's activity expresses this plan and not the recent sensory stimulation.

Quantitative analysis

Most of the LIP neurons we studied showed the response pattern illustrated in Figs. 5 and 6. We compared the activity in the M1 period within the pairs of classes 1-5 and 3-4. Within each class pair the RF stimulus appeared at identical eccentricities (20 deg. in classes 1 and 5; 10 deg. in classes 3 and 4; see Fig. 2), but the first saccade was in the PD in only one class in each pair (classes 1 and 3). The motor plan hypothesis predicts that responses in class 4 should be absent or smaller than in class 3, and that responses in class 5 should be smaller than in class 1 (by response we mean a significant change in activity relative to the background firing rate). The memory/attention hypothesis, on the other hand, predicts no difference in M1 response between classes 3 and 4 and between classes 1 and 5.

44 neurons in our sample (31 excitatory, 13 inhibitory) with significant responses ($p<0.05$) in at least one of the congruent conditions (classes 1 and 3) had no significant M1 responses (or significantly smaller responses) in the corresponding incongruent conditions (classes 5 and 4, respectively; $p<0.05$). The activity of these neurons thus fits the quantitative predictions of the motor plan hypothesis. For several cells we could compare only one pair of classes (1 vs. 5 or 3 vs. 4) because we tested only a subset of cells (n=21) in class 5, and because some cells did not respond significantly to stimuli at 10 deg. eccentricities, used in classes 3 and 4 (making a comparison between these classes inappropriate). The overall trend of these cells, however, was that any M1 response was greater when the planned saccade was in their PD than when it was in a different direction.

This trend is apparent in a plot of the average M1 activity index values in the 5 DDS classes for these 44 neurons (Fig. 6). The activity across this population is high in classes 1 and 3 and lower or absent in

classes 4 and 5. Analysis of variance revealed a significant effect of DDS class on the M1 index value (ANOVA with 4 d.f., $F=13.59$, $p<0.001$). Post-hoc tests showed significantly smaller index values in classes 2, 4, and 5 than in classes 1 and 3 (Tukey multiple comparison tests, $p<0.001$). Moreover, there was no significant difference in the population's responses in classes 4 and 5 (in which the RF was stimulated) from its response in class 2 (in which the RF was not stimulated). These response patterns are the ones predicted by the motor plan hypothesis. Thus, even though some neurons in this group show significant differences in only one of the two pairs of classes tested (class 1 vs. 5 and class 3 vs. 4), the neurons as a group carry a signal that unequivocally encodes the next planned saccade.

The remaining 8 neurons (7 excitatory, 1 inhibitory) had activity consistent with the memory/attention hypothesis. Their M1 responses in classes 4 and 5 were not significantly lower than in classes 3 and 1, respectively ($p>0.05$). The same pattern of responses is evident in the plots of these neurons' activity index values (Fig. 7). These neurons' activity did not contain, therefore, any information specific to the saccade plan.

This group of neurons did not encode the planned saccade at the population level either. Analysis of variance revealed only a trend toward a significant effect of DDS class on the M1 index value (d.f.=4, $F=1.993$, $p=0.139$). This was due to a trend of class 2 index values towards being lower than in all the other classes. When considered individually, the population response in class 2 was significantly smaller than the population responses in classes 1, 3, and 4 (t test, $p<0.05$; class 5 could not be compared because only 1 of these 8 neurons was tested in this class). No significant differences were observed among classes 1, 3, 4, and 5 (Tukey multiple comparison test, $p<0.05$). Such a response pattern is predicted by the sensory memory hypothesis.

Responses occurring between the first and second saccades

The two hypotheses of our study also predict different response patterns in the M2 period (between the first and second saccades; Fig. 1). Quantitative analysis of this activity cannot produce results as specific as for the M1 activity because the intersaccadic intervals were often short (100-200 ms) and of variable duration. Memory/planning responses in this period could thus be contaminated by postsaccadic activity following S1 and presaccadic activity preceding S2. In spite of these caveats, most of the 44 neurons that fit the motor plan hypothesis had M2 responses in classes 2 and 4, as predicted by the same hypothesis. In class 2 these responses were present even though no stimulus had ever appeared in the neuron's RF. These responses reached significance in class 2 for 25 of the 42 (60%) neurons with motor planning activity that were tested in this class ($p<0.05$).

Of the 8 cells whose M activity reflected the stimulus' location, 6 were tested in class 2. Of these, 3 had responses in the M2 period of class 2 ($p<0.05$).

M2 activity could also be due to eye position signals. The activity of many LIP neurons is modulated by eye position so that the background firing rate increases monotonically as gaze shifts in one direction (Andersen et al., 1990b). After the first saccade the eye is at a new position, which might be reflected in a new firing rate level. In general, however, the activity of a given neuron increases as gaze moves in the direction, relative to the fovea, of the neuron's RF (Andersen et al., 1990b). We confirmed this finding for a few neurons by recording their activity while the monkey maintained steady fixation on one of 9 possible gaze position (chosen from a square array with 20 degree spacing). M2 activity in classes 2 and 4 should therefore be even *lower* than in M1 activity if it were due to such modulation by eye position, whereas we usually observed higher values.

Receptive fields vs. tuning curves

The firing of LIP neurons for different stimulus locations and saccade vectors is usually better described by broad tuning curves than by sharply defined receptive fields (Barash et al., 1991b). These neurons are thus not simply either fully active or completely silent: their firing rate changes in a graded manner with changes in the location of a visual stimulus and in the amplitude and direction of the saccade required to foveate it. We would therefore expect that the firing rate during the memory period of the DDS task would reflect the direction and amplitude of the upcoming saccade according to a neuron's very broad tuning curve.

Neural activity in DDS class 5 confirmed this prediction for different saccade directions. The M activity before S1 and between S1 and S2 in class 5 generally matched the activity observed in single memory saccade trials for those directions. A clear example is offered by a cell that was excited before saccades into one quadrant and inhibited before saccades into the adjacent quadrant (Fig. 8). This cell produced clear excitatory M activity for single memory saccades directed up-left, up, and up-right (Fig. 8a-b), and clear inhibitory M activity for down-left saccades (Fig. 8c). We tested this neuron in a version of class 5 with T2 falling in its excitatory RF and S1 in its inhibitory RF. The neuron showed *inhibitory* M1 activity (Fig. 8d), again reflecting the planned saccade and not the recent sensory stimulation. The motor planning activity of each LIP neuron thus contributes to the next saccade plan in a graded manner over the entire range (360 deg.) of possible directions.

Interaction of sensory response and motor plan

The neuron in Fig. 8 shows an interesting pattern of responses to sensory stimuli. A stimulus in the left upper quadrant, when presented alone (Fig. 8a), elicits a strong response during the stimulus-presentation period itself (LS response). When the same stimulus (T2, Fig. 8d), however, appears following an inhibitory stimulus in the

lower left quadrant, the excitatory response is absent. We observed absent or attenuated LS responses, following a stimulus outside the excitatory RF, in 10 neurons. It is as if a neuron decided, based on the first stimulus, the direction of the saccade to be made, and from then on maintained the appropriate firing rate for that saccade plan, thus attenuating or abolishing any sensory responses to subsequent stimuli. Some neurons thus appear to express the saccade plan quite early, and to respond to sensory stimuli based on the context of the current saccade plan.

Timing of onset of M2 responses

During the delay period of classes 2 and 4 the monkey is planning a saccade away from a given neuron's RF, but also knows that immediately after the first saccade he will make the second one towards its RF. The second motor plan is reflected in the M2 activity of the neurons encoding the next saccade. We asked when this activity started appearing relative to the first saccade. It has been reported that the RF's of some LIP neurons shift in advance of a saccade that will bring a visual stimulus into their RF, a process that has been termed "predictive remapping" of the RF (Duhamel et al., 1992a). We wanted to know whether the motor planning activity was remapped predictively in the LIP neurons we studied. To this end we measured the latency of onset of M2 responses relative to the beginning of the first saccade in class 2. In this class no stimulus ever appears in a neuron's RF, allowing us to measure the motor planning component of a neuron's activity without the contamination of stimulus-related activity.

Fig. 9 shows the latency of class 2 M2 responses for 23 neurons with clear M2 responses. These responses started well after the beginning of the saccade in most neurons (mean = 141 ms; s.e.m. = 17 ms). In only two neurons did the activity begin before saccade onset, and in all but 4 neurons the activity started later than 100 ms after saccade onset. The response field of the motor planning activity of almost all

LIP neurons in our sample thus did not remap in advance of the saccade.

DISCUSSION

Motor intention

The most important finding of this study is that the activity of most neurons in area LIP encodes the next planned saccade. It had already been shown using a double saccade task that LIP neurons become active before a saccade into their motor response field even in the absence of sensory stimulation (Gnadt et al., 1988; Barash et al., 1991b). In those experiments, however, the monkeys made the saccades immediately after the stimuli had appeared. The components of neural activity related to sensory memory and motor planning were thus potentially confounded with those related to sensory stimulation and saccade execution. In the present study we imposed a delay between presentation of the stimuli and saccade execution. During the delay the monkeys had to hold in memory two locations and plan the next saccade while maintaining fixation. The M activity we recorded could thus only be related to covert processes such as sensory memory, allocation of visual attention, and saccade planning.

We created a conflict, in our experiments, between the location of sensory stimulation and the planned saccade by presenting a stimulus (the second target), in the neurons' sensory response field (the RF), that was not the goal of the next saccade. M activity in this condition was reduced or completely absent in most neurons compared to when the same stimulus was also the goal of the next saccade. Sensory stimulation alone is thus not sufficient to activate most LIP neurons during the delay period. We also replicated the finding of Gnadt et al. (1988) and Barash et al. (1991b) that LIP neurons do not require sensory stimulation to manifest motor planning activity. Most became active between two saccades if the second one was into their motor response field, even when no stimulus had appeared in their sensory response field. Thus in the period between a visual stimulus and a

saccade LIP neurons clearly encode, in motor coordinates, the next intended saccadic eye movement.

As mentioned above, if the M activity were a memory trace of a stimulus' location or an attention shift to that location, it should have been triggered every time a stimulus appeared in a neuron's RF, independently of what saccade was being planned. A few neurons in our sample showed this pattern of responses. Their M activity thus likely subserves higher-level processing of the visual world, by maintaining representation of a stimulus' location and/or allocating attention. This role is consistent with area LIP's extensive inputs from other visual areas (Blatt et al., 1990) and with the strong responses of this area's neurons to visual stimuli. Bushnell et al. (1981) have shown that parietal neurons produce enhanced responses to visual stimuli when a monkey actively attends to them. It is thus possible that some LIP neurons play a role in allocating attention in the visual field.

The large majority of the neurons we studied, however, responded as predicted by the motor plan hypothesis. This prediction is that a neuron should only be active during the memory period of a delayed double saccade trial if the next saccade were to be made in its preferred direction. Thus even if the second target were flashed in the receptive field, it should not evoke M activity. The cell should only become active *after* the first saccade, if the second saccade were in its preferred direction. About four fifths of the cells we recorded showed this response pattern.

By observing neural activity during an imposed delay between a monkey's perception of a visual stimulus and an orienting eye movement to it, we have identified a neural correlate of the animal's oculomotor intention. The M activity of most LIP neurons reveals what saccade the animal intends to make next. We want to emphasize that the pattern of responses we observed cannot be explained as only reflecting shifts of attention in the visual field. Such covert attention shifts are reliably elicited by behaviorally relevant stimuli (Posner et al., 1984), such as saccade targets. In the DDS task the monkeys had

to attend to two spatial locations, as both were targets of future saccades. The M activity of most neurons, however, was maintained only if the stimulus in the RF was also the target of the next saccade, and thus did not reflect the allocation of visual attention.

It may be argued that a saccadic eye movement is preceded by a special type of attention shift, characterized by the fact that it *always* precedes a saccade, and that the M activity of LIP neurons reflects this particular type of attention rather than oculomotor intention. These two terms—an attention shift to a location that is contingent upon a planned saccade to the same location, vs. the saccade plan itself—describe, however, operationally equivalent processes that cannot be distinguished experimentally.

Our results show that the M activity of most LIP neurons is not solely determined by sensory and perceptual processes but also carries information about the upcoming eye movement. A related question is whether the activity predicts the upcoming saccade in an obligatory fashion. LIP neurons often discharge just before, during, and after a saccade (Barash et al., 1991a). It is thus possible that the M activity is simply an early expression of the activity that will coincide with the saccade, that is, the motor command rather than the motor plan. If the M activity truly reflects the intention to make the saccade, on the other hand, it should be independent of the actual execution of the movement. We performed another set of experiments to address this question. These are described in the following chapter. Briefly, we found that the M activity is indeed not contingent on actual saccade execution, and thus encodes the animal's intention rather than purely premotor processes.

Timing of onset of motor planning activity

Duhamel et al. (1992a) have found that, when a saccade is about to move a visual stimulus into the RF of an LIP neuron, the visual response of a group of LIP neurons anticipates the saccade. The authors suggested that the response fields of these neurons temporarily shift, near the time of saccade execution, in the direction of

the upcoming saccade; that is, they are predictively remapped. We found that no such anticipation occurs for the motor planning activity of LIP neurons. This plan is made manifest in the neurons' activity only when the upcoming saccade is truly the one in a given neuron's motor planning field. As Fig. 9 shows, most LIP neurons remain silent as long as an intervening saccade outside their motor field is being executed. The motor planning response fields of these neurons are not predictively remapped in advance of saccades, and thus never shift. Their activity is an unambiguous code for the saccade being planned at the moment.

We do not know whether the sensory responses of the LIP neurons in our study were predictively remapped. We could not ask this question in the paradigm of this experiment, because the sensory responses were extinguished by the time the first saccade was made at the end of the delay period.

Context-dependent visual responses

Does a stimulus, if it falls in the receptive field of the cell, always evoke a visual response (activity during the stimulus presentation period)? Although this appears to be true for the majority of LIP cells, in some the sensory response is attenuated or absent if the stimulus is the second in the sequence. We suggest that this is a result of the behavioral significance of the stimulus: when presented first it cues the next saccade, whereas when second it cues the second saccade. This suggests that the behavioral significance of the stimulus in part determines the "visual" response to it of some LIP neurons.

These findings are reminiscent of the observations of many workers in other higher order motor regions of the brain, where "sensory" responses to stimuli are typically only observed when the stimuli are cues to move. For instance, Godschalk et al. (1985) observed visual responses in the premotor cortex only if the stimuli cued an arm movement. Similarly, Seal and Commenges (1985) recorded responses to auditory stimuli in area 5 of the PPC only when these stimuli instructed the monkey to reach for a target.

Role of area LIP activity

It has been suggested (Goldberg et al., 1989, 1990; Duhamel et al., 1992a; Goldberg and Colby, 1992) that a major role of area LIP is the spatial analysis of the visual scene. The hypothesis is that one function of area LIP is to remap the visual scene before a saccade is made in order to predict what the reafferent visual input will be after the eye movement. Such a role would be consistent with the "where" function (i.e. the localization of visual stimuli) ascribed to the occipito-parietal pathway of cortical visual areas (Ungerleider and Mishkin, 1982), which includes area LIP. By spatially dissociating sensory stimulation from the planned saccade, however, we have shown that most LIP neurons encode—during an imposed delay—the monkey's intention to make the next saccade, independently of sensory stimulation. Thus, although the locations of sensory stimuli is clearly encoded in the responses of LIP neurons—in their stimulus-period responses, as well as in the memory-period activity of a few LIP neurons—, these neurons also carry an unambiguous code of the monkey's oculomotor intention. The role of area LIP thus cannot be limited to visual spatial analysis, but extends to a participation in movement planning.

The motor planning activity of LIP neurons demonstrated in this study activity is the first expression of motor intention identified in the visual cortex. Activity playing an analogous role has been recorded in another sensory association area, area 5, where a population of neurons encodes the upcoming arm movement (Kalaska and Crammond, 1990). Activity encoding the upcoming movement is also common in cortical areas in the frontal lobe that are directly involved in movement preparation and execution, such as the motor (Georgopoulos et al., 1989), premotor (Bruce et al., 1985; Weinrich and Wise, 1982; Wise et al., 1992), and prefrontal cortex (Funahashi et al., 1990).

The broad directional tuning of LIP neurons' response fields (Barash et al., 1992b) is consistent with a role of this area in encoding

the next planned saccade. Such tuning allows an entire population of neurons to participate in the coding of one vector, creating a distributed representation of a single signal which is much less sensitive to the noise in the firing rate of individual neurons than a code expressed only at the single cell level. This coding strategy seems commonly employed in nervous system structures encoding movements (Georgopoulos et al., 1986, 1988; Sparks et al., 1988). The drawback of broad tuning curves is that only a few spatial signals can be encoded simultaneously. This feature of LIP neurons thus also makes it less likely that area LIP represents the visual scene, as has been suggested (Duhamel et al. 1992a).

Our results, combined with the fact that area LIP is indeed located in the "where" cortical visual pathway (Ungerleider and Mishkin, 1982), support a revision of the role of the occipito-parietal pathway recently proposed by Goodale and Milner (1992). These authors described a patient with a posterior parietal lesion who could visually discriminate the sizes of various objects but could not adjust the size of her hand grip when grasping the objects (Goodale and Milner, 1992). A similar deficit was obtained in monkeys following ablation of the cortex in the superior parietal lobule (Halsband and Passingham, 1982). Goodale and Milner (1992) proposed that the role of the occipito-parietal cortical stream is to figure out not only "where" a sensory stimulus is, but also "how" to prepare an action based on the stimulus' spatial attributes. Our results would place the memory signals of LIP neurons near the final answer to such a question. Given the location of a sensory stimulus, LIP neurons "prepare" an intended saccade that will align the stimulus' location with the fovea.

By expressing only an intention to move the eyes, LIP neurons encode the most primal form of a motor plan, one that unambiguously encodes the action being prepared but which can be freely altered as new behavioral contingencies are perceived. Area LIP thus seems to establish a direct interface, along the extrastriate visual pathways, between the sensory and motor domains by encoding the transition from spatial perception to movement plan in the activity of its neurons.

FIGURE LEGENDS

Figure 1:

Timing of stimuli and eye movements in the DDS task. Horizontal bars indicate the appearance of the visual stimuli. Below these traces is a sketch of a typical eye position trace illustrating the monkey's behavior in the task. T1 = first visual target; T2 = second visual target; M1 = memory, or delay, period; S1 = first saccade; S2 = second saccade.

Figure 2:

Spatial arrangement of stimuli in the 5 classes of the DDS task and predictions of M1 period responses. Each panel in the top row shows the location of the fixation point (cross), a sample neuron's receptive field (dashed semicircle), the locations where the two visual targets appear (dots), and the amplitude and direction of the saccade the monkey must make (arrows). The table below the row of diagrams indicates which stimuli, in each class, fall in this neuron's receptive field and which saccades are in the neuron's preferred direction. The table's bottom two rows whether or not M1-period activity is expected in each class according to the memory/attention hypothesis and according to the motor plan hypothesis. FP = fixation point; T1, T2 = first and second visual targets, respectively; RF = receptive field; PD = preferred direction.

Figure 3:

Activity of an area LIP neuron in two types of single memory saccade trials. The abscissa in each panel represents time (100 ms/division) during each trial of the task. Within each panel are plotted, from top to bottom: rasters of tick marks representing the occurrences of action potentials in each trial (1 trial/row); a time histogram (bin width = 50 ms) of the neuron's average rate of action potential firing over all trials (20 Hz/division); and a trace of the monkey's vertical eye position (30

deg./division). Onset and offset times of stimuli during the trials are indicated both by the thin vertical lines within each panel and by the thick horizontal lines below each panel. To the left of each panel is a sketch of the arrangement of the visual target (dot) and the saccade made by the monkey (arrow) relative to the neuron's receptive field (dotted semicircle). FP, fixation point; RF, neuron's receptive field; T, visual target; M, memory, or delay, period. *A*) The visual target falls in the neuron's RF and the saccade is in the neuron's PD. *B*) The visual target falls outside the RF and the saccade is in the direction opposite the PD.

Figure 4:

Activity of an LIP neuron in classes 1-4 of the DDS task. As in Fig. 3 each panel has a plot that includes, from top to bottom, the spike rasters for each trial, the time histogram (bin width 50 ms) of the firing rate (20 Hz/division in *A-C*, 25 Hz/division in *D-E*), and the horizontal and vertical eye positions (30 deg./division) (abscissa: 100 ms/division). The vertical dotted lines and the thick horizontal lines below each panel again show the onset and offset of the visual stimuli. The diagrams to the left of each panel show the spatial arrangement of the first and second target (T1 and T2, respectively), the first and second saccades (arrows), and the neuron's receptive field (RF). *A*) Class 1; *B*) class 2; *C*) class 3; *D*) class 4.

Figure 5:

Activity of an LIP neuron in a single memory saccade (*A*) and in class 5 of the DDS task (*B*). The panels are arranged as in Figs. 4 and 5, using the same abbreviations. The vertical scale for the histogram is 10 Hz/division.

Figure 6:

Mean values of the M1-period activity index, I_a, in the 5 DDS classes for the 44 LIP neurons whose M1 responses fit the motor plan

hypothesis by statistical tests. Shown for each class are the class mean ± standard error. The number of index values is not the same in each class (see methods). These numbers are n = 33 (classes 1 and 2), n = 40 (classes 3 and 4), and n = 20 (class 5).

Figure 7:

Mean values of the M1-period activity index, I_a, in the 5 DDS classes for the 8 LIP neurons whose M1 responses fit the memory/attention hypothesis by statistical tests. Shown for each class are the class mean ± standard error. The number of index values is not the same in each class (see Methods). These numbers are n = 3 (classes 1 and 2), n = 8 (classes 3 and 4), and n = 1 (class 5).

Figure 8:

Tuning of M activity for the direction of the next saccade. Panels *A-C* show the activity of an LIP neuron in single memory saccade trials. Panel D shows the activity of the same neuron in class 5 of the DDS task.

Figure 9:

Latencies of the M2 responses in class 2 of 23 LIP neurons with motor planning memory activity. The latency indicates the onset of these responses relative to the beginning of the first saccade. Each dot in this plot marks the latency of one neuron.

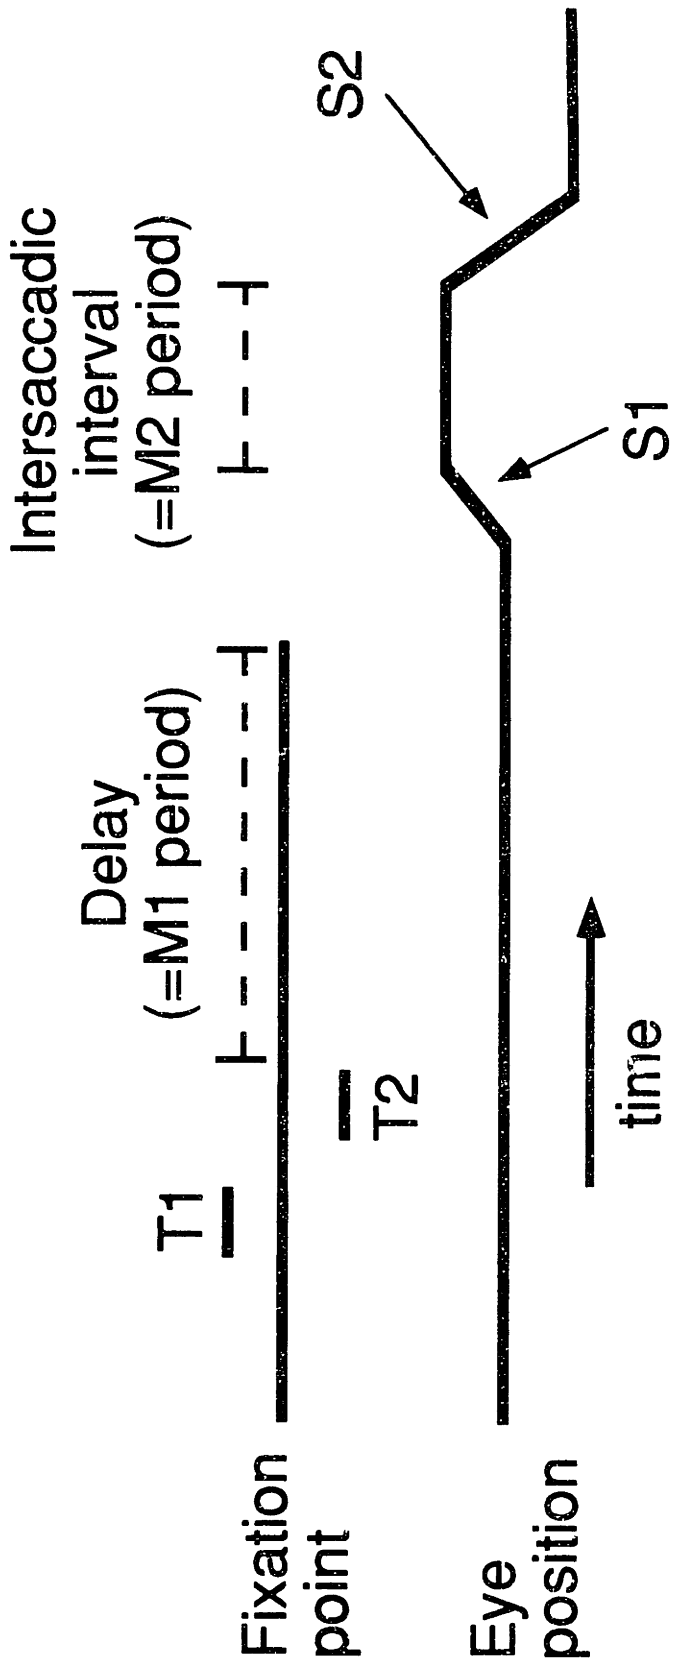
Intersaccadic interval (=M2 period)
S2
S1
Delay (=M1 period)
T1
T2
time
Fixation point
Eye position

Figure 1

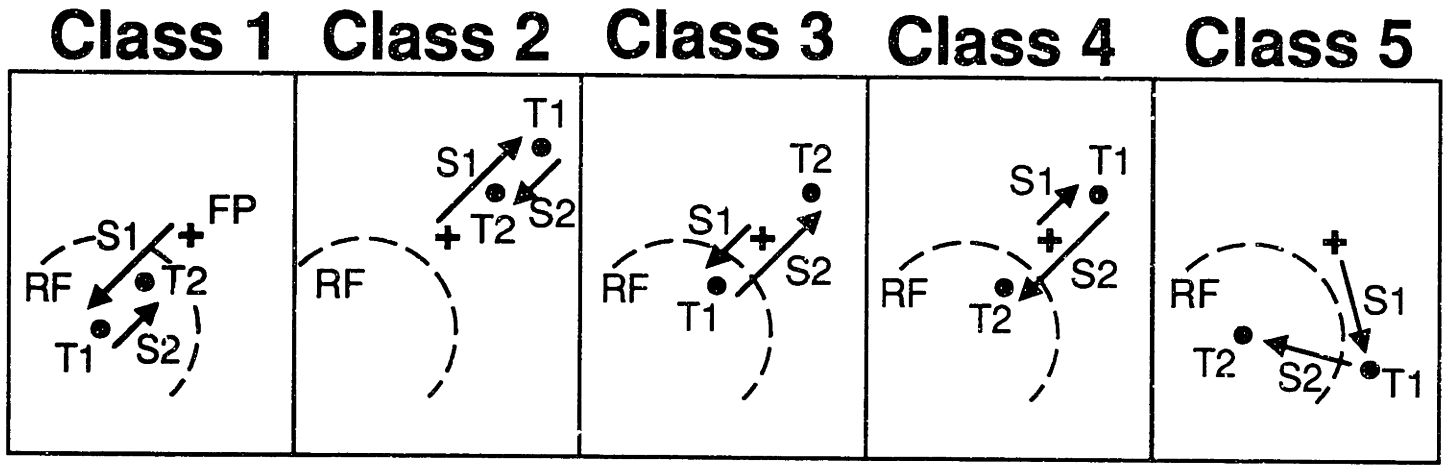

	Class 1	Class 2	Class 3	Class 4	Class 5
T1:	RF		RF		
T2:	RF			RF	RF
S1:	PD		PD		
S2:		PD		PD	
M1: memory/ attention	+	—	+	+	+
M1: motor plan	+	—	+	—	—

Figure 2

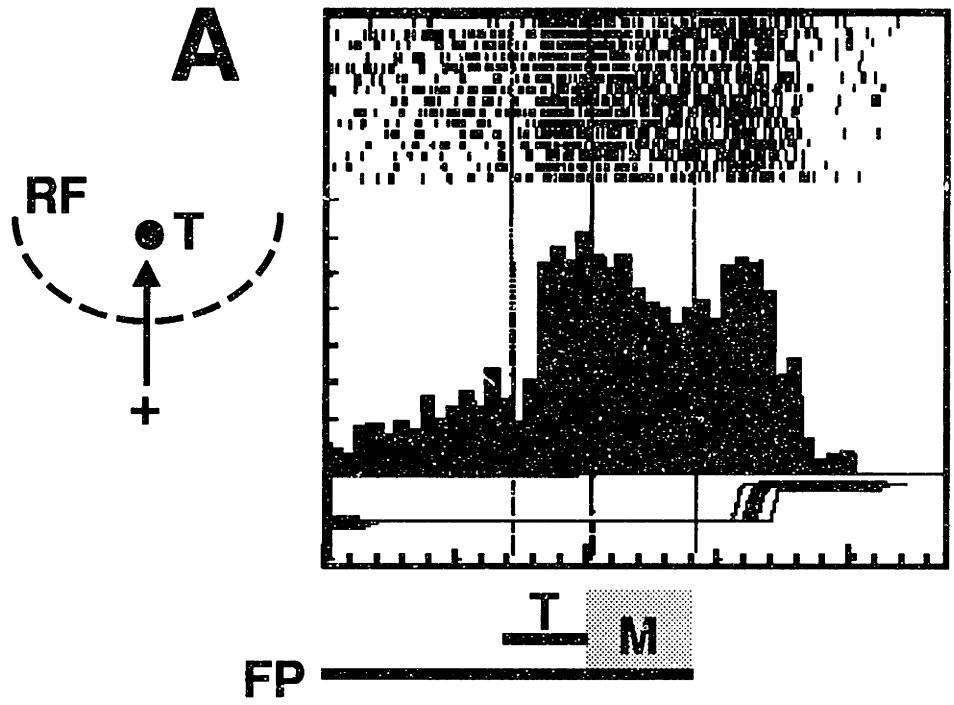

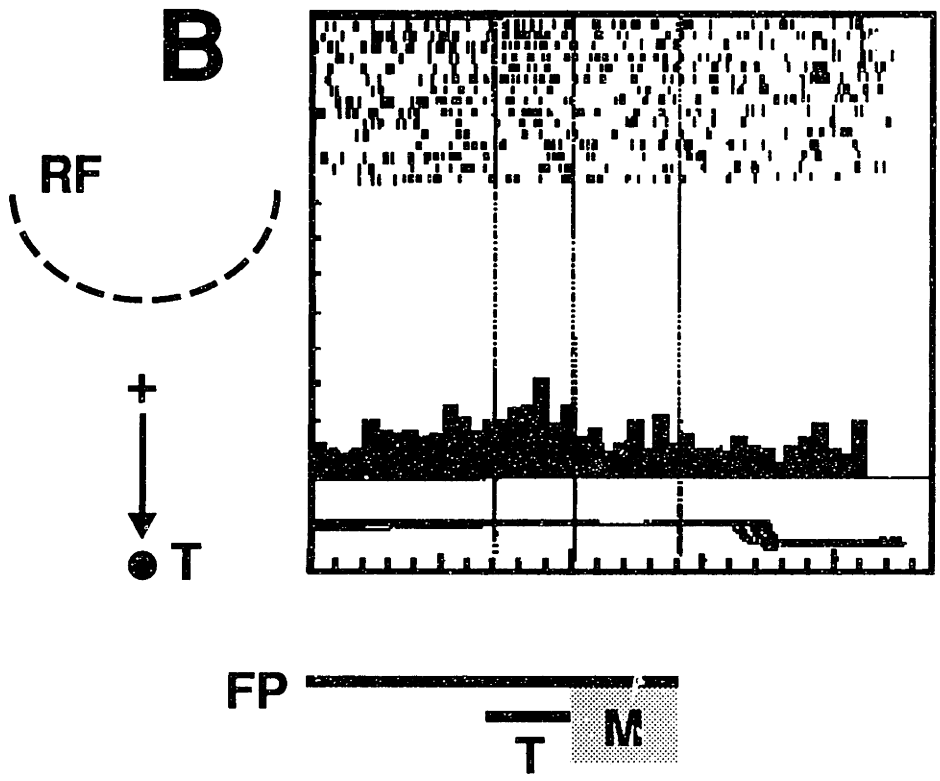

Figure 3

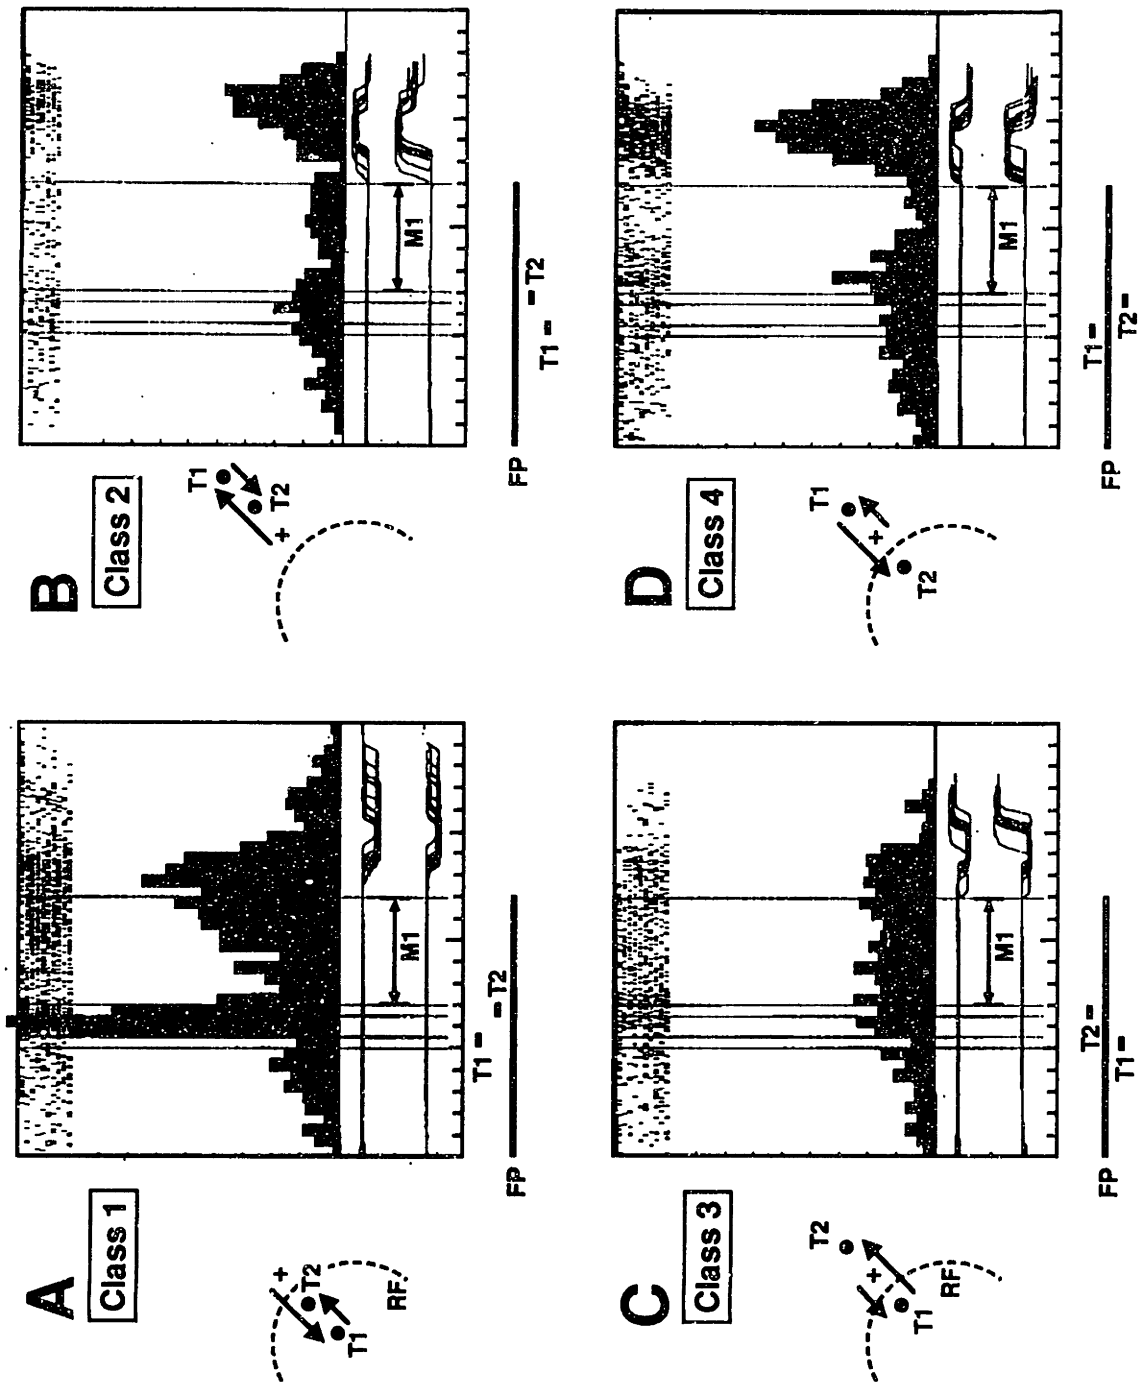
A
Class 1
T1
T2
RF
M1
FP
T1 - -T2
B
Class 2
T1
T2
M1
FP
T1 - -T2
C
Class 3
T2
T1
RF
M1
FP
T2-
T1-
D
Class 4
T1
T2
M1
FP
T1-
T2-

Figure 4

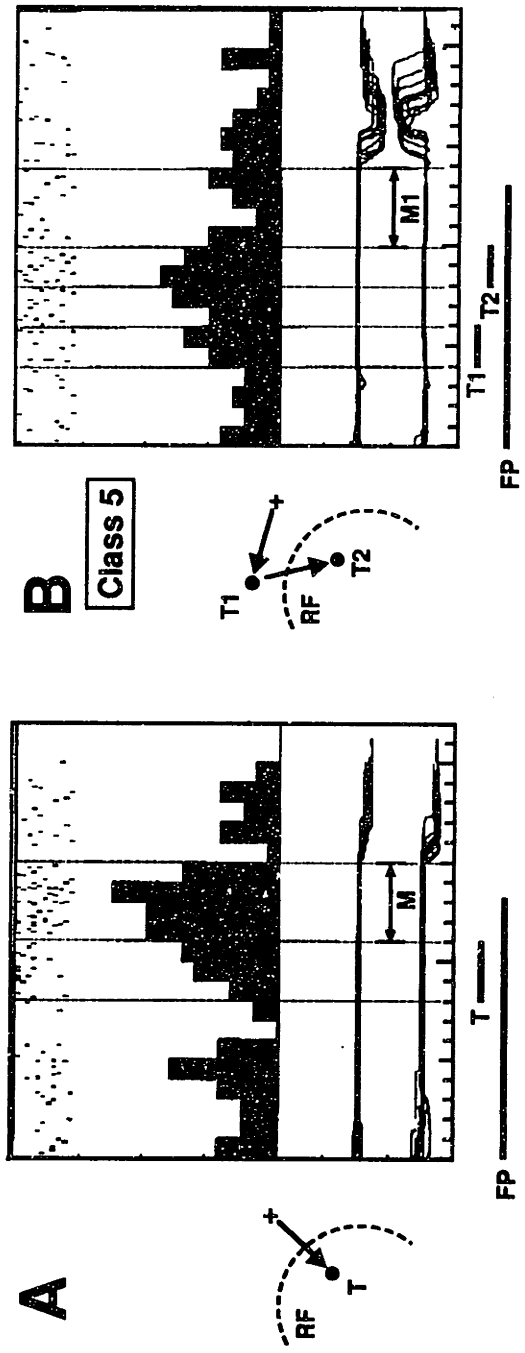

Figure 5

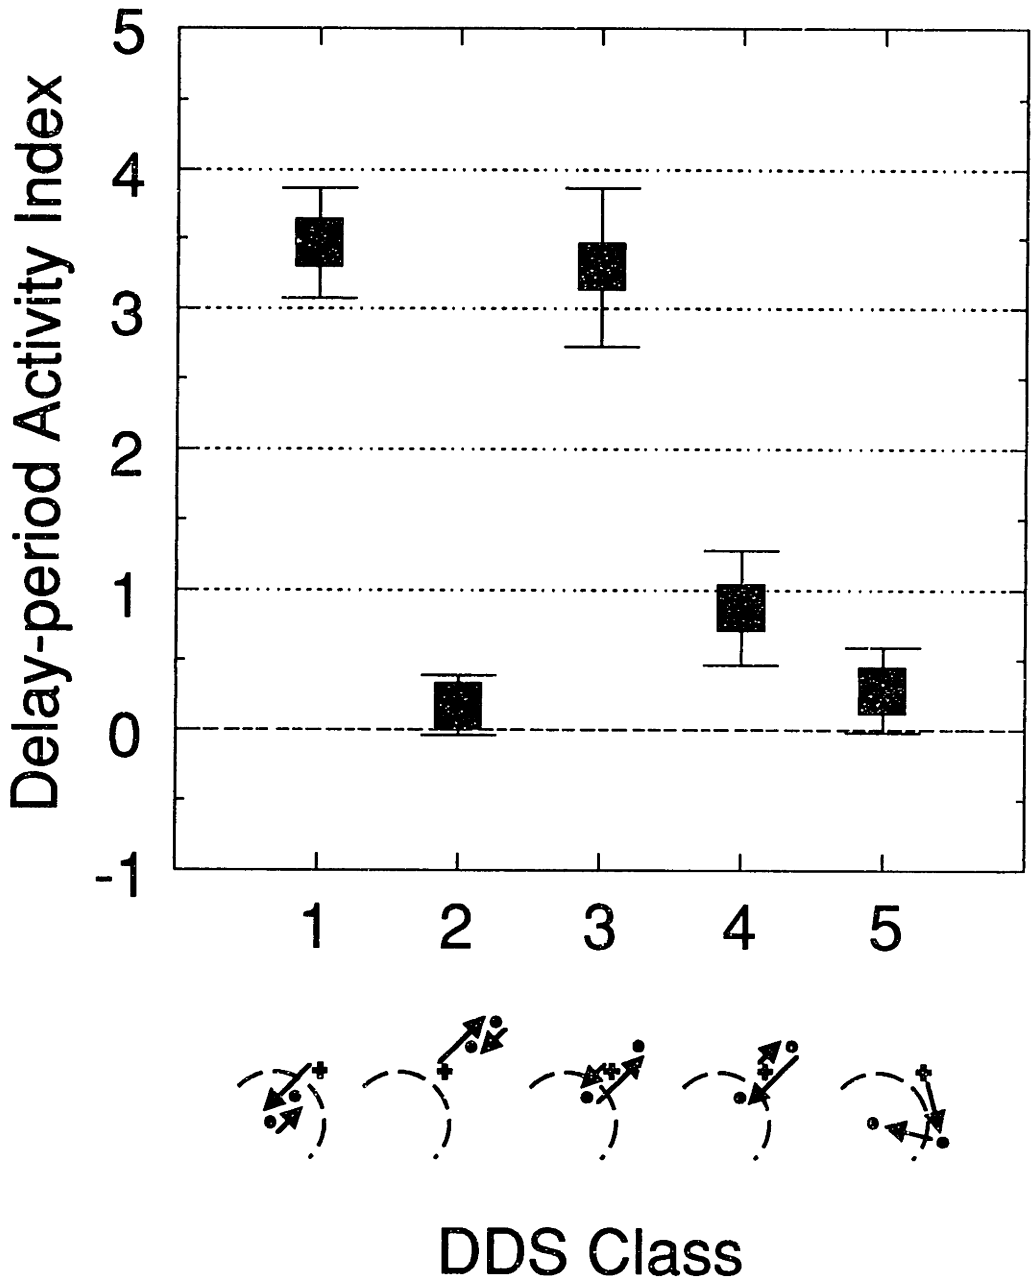

Figure 6

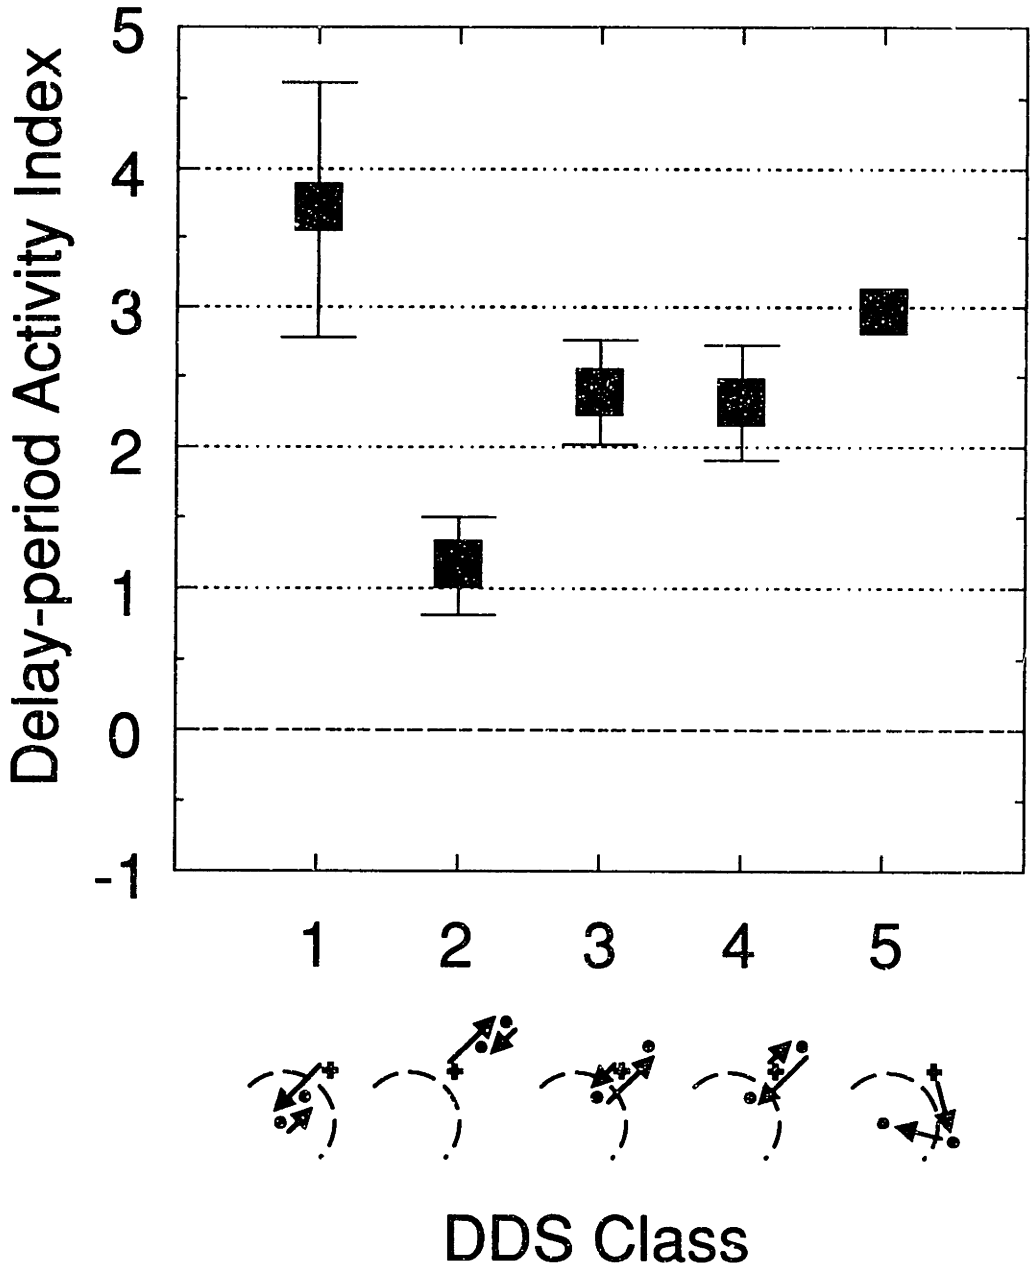

Figure 7

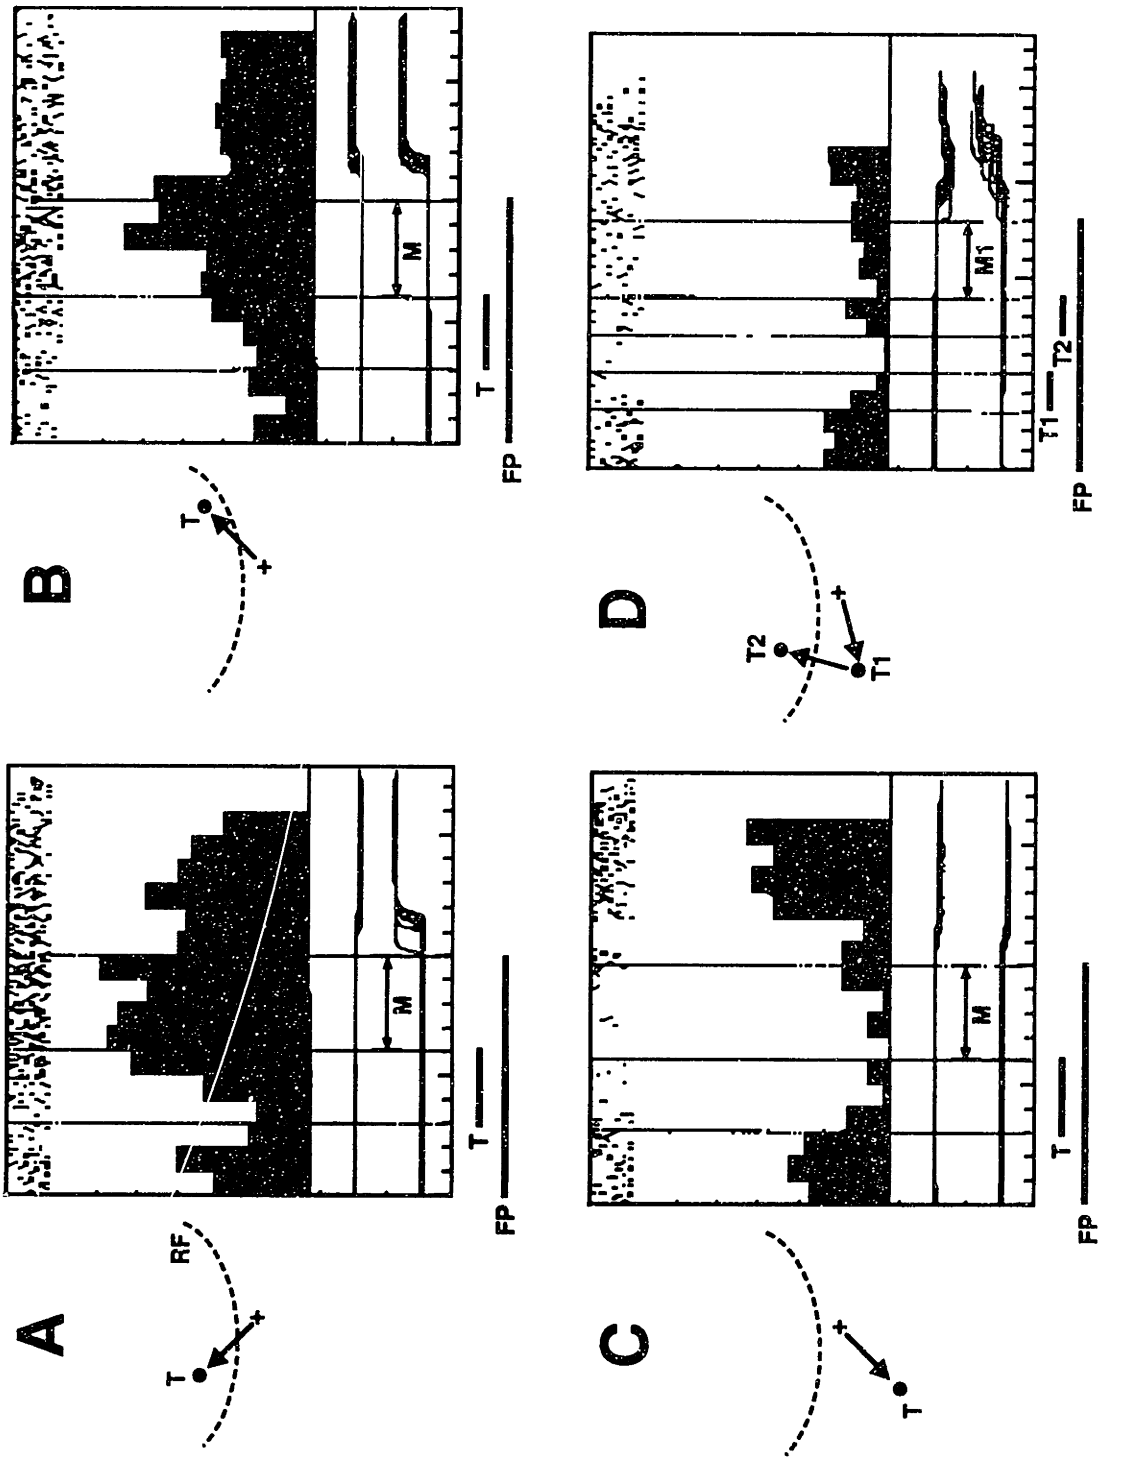

Figure 8

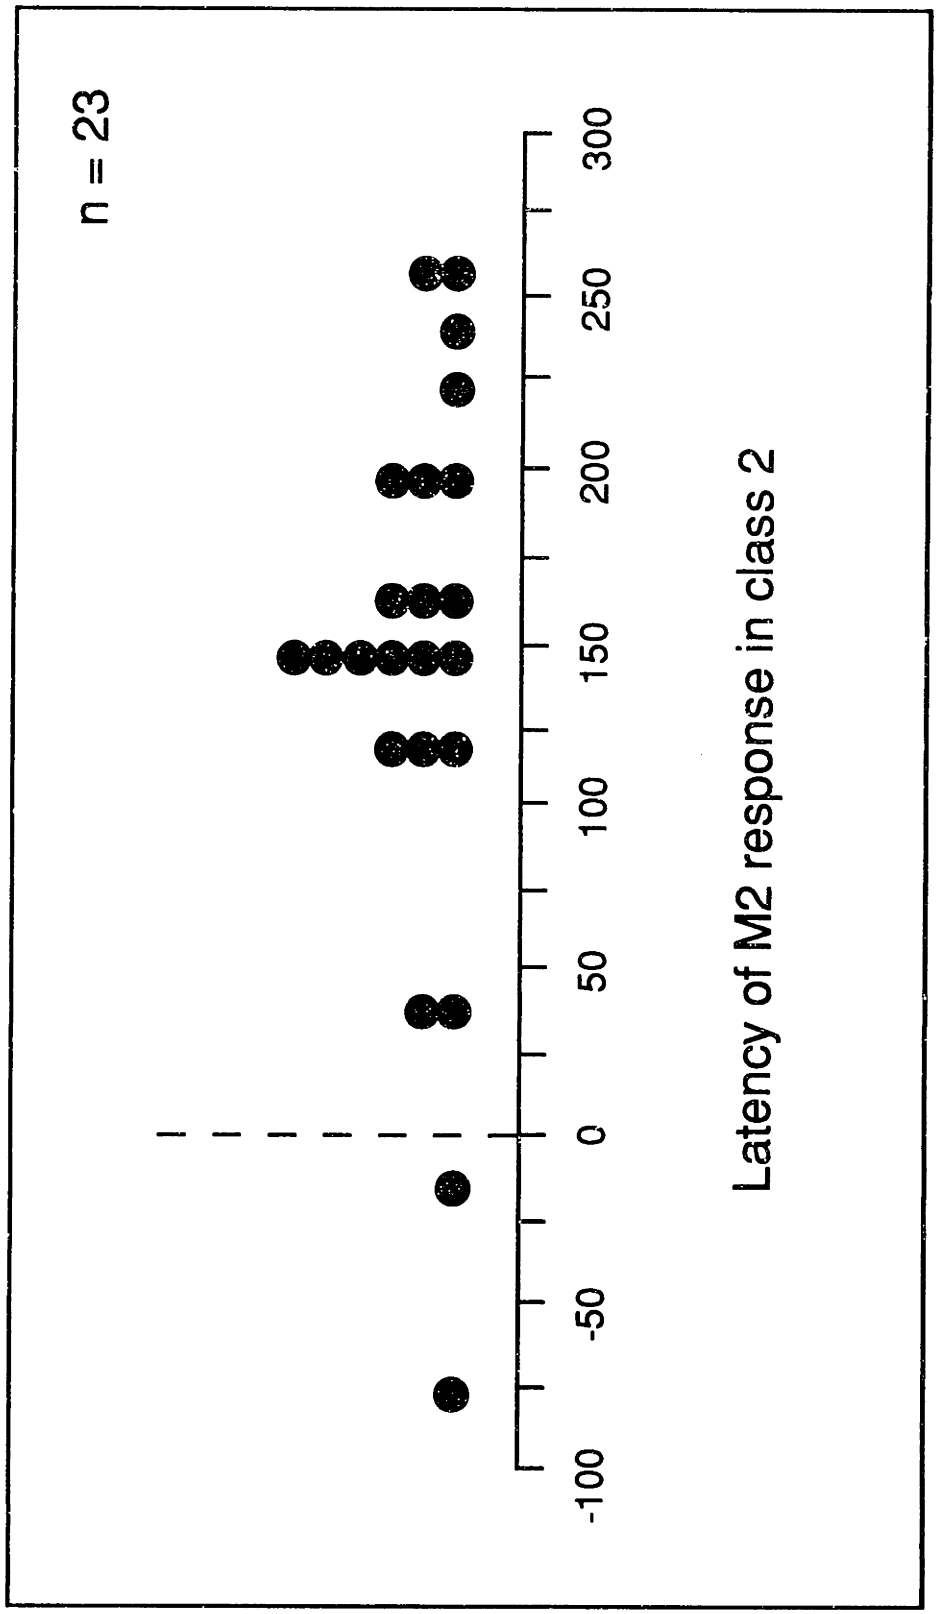
n = 23
-100
-50
0
50
100
150
200
250
300
Latency of M2 response in class 2

Figure 9

Chapter 8

Motor intention activity in area LIP. II. Changes of motor plan

SUMMARY

In the previous chapter we reported that most neurons in the lateral intraparietal area (area LIP) of the monkey's posterior parietal cortex (PPC) encode the next intended saccadic eye movement. This result predicts that, should the monkey change his intention of what the next saccade will be, LIP activity should change accordingly to reflect the new plan. We tested this prediction by training a monkey to change his saccadic plan on command and recording the activity of LIP neurons across plan changes.

We trained rhesus monkeys (*Macaca mulatta*) to maintain fixation on a light spot as long as this spot remained on. During this period we briefly presented one, two, or three peripheral visual stimuli in sequence, each followed by a delay (memory period, M). After the final delay the fixation spot was extinguished and the monkey had to quickly make a saccade to the location of the last target to have appeared. The monkey could not predict which stimuli, nor how many, would appear on each trial. He thus had to plan a saccade to each stimulus as it appeared and change his saccade plan whenever a stimulus appeared at a different location.

We recorded the M period activity of 81 area LIP neurons (from 3 hemispheres of 2 monkeys) in this task. We predicted that, if a neuron's activity reflected the monkey's planned saccade, its activity should be high while the monkey planned a saccade in the neuron's preferred direction (PD), and low while the planned saccade was in the opposite direction. The activity of most of the neurons in our sample changed in accordance with our hypothesis as the monkey's planned saccade changed.

In one condition the monkey was instructed by visual stimuli to change his plan from a saccade in the neuron's preferred direction to a saccade planned in the opposite direction. In this condition activity decreased significantly ($p<0.05$) in 65 (80%) out of 81 neurons tested. These neurons' activity changed to reflect the new saccade plan even though the cue for this change was not presented in their RF.

As a control we randomly interleaved, among trials requiring a plan change, trials in which the monkey had to formulate two consecutive plans to make a saccade in the neuron's preferred direction. The activity remained unchanged ($p<0.05$) in 22 out of 31 neurons tested (79%), and thus remained consistent with the maintenance of the same saccade plan.

In a variant of the task the cue to the location of the required saccade was either a light spot or a noise burst from a loudspeaker. Of 22 neurons tested in this task, 16 (73%) showed activity changes consistent with plan changes cued by visual or auditory stimuli.

Alterations in the monkey's intentions, even in the absence of overt behavior, are manifested in altered LIP activity. These activity changes could be induced whether visual or auditory cues were used to indicate the required plan changes. Most LIP neurons thus do not encode only the locations of visual stimuli, but also the intention to direct gaze to specific locations, independently of whether a gaze shift actually occurs.

INTRODUCTION

We continue in this chapter our investigation of the role that neurons in the monkey's lateral intraparietal area (area LIP) play in the preparation of saccadic eye movements. Many neurons in this area of the posterior parietal cortex (PPC) are active from the moment a stimulus appear in their visual receptive field (RF) until a saccade to foveate that stimulus is completed (Gnadt and Andersen, 1988). If a delay is imposed between stimulus presentation and saccade execution (memory saccade task; Hikosaka and Wurtz, 1983), these neurons show distinct stimulus-related and saccade-related (SR) discharges, as well as elevated activity during the delay ("memory"-period, or M, activity) (Gnadt & Andersen, 1988; Barash et al., 1991a). M and SR responses are especially prominent in area LIP (Barash et al., 1991a, b), an area distinguished from other PPC areas by its strong connections with other saccade-related regions, such as the superior colliculus and frontal eye fields (Lynch et al., 1985; Andersen et al., 1990a; Blatt et al., 1990).

During the delay in a memory saccade task the monkey receives no sensory input and makes no eye movement, but must remember where the stimulus has just appeared and plan an appropriate saccade to that location. We have been studying the M activity to see what role it plays in the sensorimotor transformations related to saccadic eye movements.

On the basis of double saccade and back saccade tasks (Gnadt and Andersen, 1988; Barash et al., 1991b), we have demonstrated that LIP cells code in motor coordinates: they become active if a saccade is planned into their motor field, even it no visual target ever falls within their receptive field. Moreover, these results and the results of the delayed double saccade experiments described in the accompanying study (previous chapter) suggest that the sustained (M) activity of in LIP cells reflects the monkey's intention or motor plan to make the next saccade into its motor field. The M activity is not a sensory memory; it reflects the vector of the forthcoming saccade.

We reasoned that if this activity did indeed reflect the monkey's motor plan, then alterations of this plan, even in the absence of overt behavior, should be manifest in altered LIP activity. We report here the results of experiments using a "change of motor plan" (CP) paradigm designed to test this hypothesis. In this paradigm, one, two or three targets were presented sequentially during the fixation period. The monkey did not know how many, nor which, targets would be presented, as the different trial types were pseudorandomly interleaved. He presumably planned to saccade to the first target when it appeared, and then changed his plan if a subsequent, different, target appeared. To insure that the monkey planned to make the eye movement during the delay period and not after the fixation light was extinguished, he was required to saccade within a very short reaction time in order to receive a reward. We were thus able to correlate changes in the motor plan with alterations in single unit activity in LIP. The results support our conjecture that neuronal activity in LIP reflects the monkey's intention to make the next saccade.

METHODS

The methods we used in this experiment are the same as those of the companion study (described in the previous chapter), except for the behavioral tasks and data analysis. We describe these next.

Behavioral tasks

Each monkey learned to perform several tasks involving saccades for the purposes of several studies. The ones used in this study are the visual memory saccade task, the visual change of plan task, and the visual-auditory change of plan task.

A *memory saccade (MS)* trial started when a spot was turned on directly in front of the monkey and the monkey started fixating on it. After 800 ms of fixation a peripheral stimulus was presented for 300 ms. The monkey had to continue fixating for another 400 ms after stimulus offset (delay, or memory, period). At this point the fixation

spot was extinguished and the monkey had to make a saccade, in the dark, to the remembered location where the stimulus had appeared. The stimulus was placed at an eccentricity of 5 to 25 degrees along one of 8 directions (the 4 cardinal and the 4 diagonal directions).

The responses of most LIP neurons during a MS trial consist of at least one of three components. These are a sensory response (LS, appearing during light stimulus presentation), a saccade-related response (SR, coincident with saccade execution), and sustained activity during the delay between stimulus presentation and fixation spot offset (memory activity, M) (Gnadt et al., 1988; Barash et al., 1991a, b). These signals are spatially tuned. LS responses vary across the visual field, being strongest for stimuli in a circumscribed sensory response field (receptive field, RF). SR responses vary with saccade direction (and to some extent with amplitude), reaching a maximum for saccades in the neuron's preferred direction (PD). The spatial tuning of the visual and saccade-related responses in LIP generally coincide, that is, the RF is in the same direction, relative to the fovea, as a saccade in the neuron's PD. The spatial tuning of the M activity generally matches that of the LS and SR responses (Barash et al., 1991b). LIP neurons thus have up to three spatially tuned response fields (sensory, memory, and motor), which are generally aligned. We used MS trials to locate each neuron's visual RF and saccadic PD.

If a neuron had sustained M activity we then tested it in the *change of motor plan (CP)* paradigm. The CP task consisted of up to 8 classes of trials, all involving stimuli at one of 2 locations. Location A (the neuron's *preferred* location) was in the RF, while location B (the nonpreferred location) was outside the RF, diametrically opposed to location A (. 1). In all classes the monkey had to maintain gaze on a fixation point (FP) while various combinations of stimuli appeared, and then, after FP offset, make a saccade to the remembered location of the last target to have appeared.

Fig. 1 shows the time course of onset and offset of the FP and the visual stimuli in each of the 8 classes. The first two classes were standard memory saccade trials with the stimulus located at A (inside

the receptive field, class 1) or B (outside the receptive field, class 2). Each trial began with the appearance of the FP, followed after 800 ms by the appearance of the visual stimulus (A or B), which stayed on for 300 ms. A delay of 400 ms (the memory, or M, period) followed the stimulus' offset. The FP was then extinguished and the monkey made a saccade, within 350 ms, to the remembered location of A or B.

Classes 3 and 4 required a change of motor plan. As in the first two classes, a stimulus (A in class 3, B in class 4) was presented for 300 ms after the monkey had been fixating for 800 ms. After the 400 ms delay period (M1), however, a second stimulus (B in class 3, A in class 4) appeared for 300 ms, and was followed by a second 400 ms delay period (M2).

We pseudorandomly interleaved trials of all classes in the CP task. The monkey could not know in advance, therefore, how many nor which targets would appear in a given trial. Because he had only 350 ms to make the required saccade after FP offset, he had to always be ready to make a saccade to the most recent target's location in case the FP was turned off. He thus presumably planned a saccade to each stimulus as it appeared and maintained that plan throughout the M period that followed. When a second stimulus appeared as in classes 3 and 4 he then had to change his saccade plan and maintain the new one during the next M period.

Classes 5 and 6 were control trials. It is possible that the mere appearance of a second visual stimulus, independent of the saccade plan, might affect the activity of LIP neurons and thus confound any activity changes related to the change in saccade plan. We controlled for this possibility by presenting two stimuli, as in classes 3 and 4, but both at the same location (A in class 5 and B in class 6).

In classes 7 and 8 we presented 3 stimuli, requiring two changes of saccade plan. In class 7 the monkey had to first plan a saccade to A, then change that plan to a saccade to B, then change it again to a saccade to A. Class 8 had a B-A-B plan sequence.

The third task used in this study is the visual-auditory change of plan (CPVA) task. This is a variant of the CP task in which the

stimulus could be either visual or auditory. We used this task to see whether any patterns of responses observed in the CP task might be specific to visual stimuli or could be elicited by nonvisual spatial cues. The visual stimulus was the same light spot used in the other tasks and was located at 8 deg. eccentricity to the right or to the left of the fixation point. The auditory stimulus was a 20-20,000 Hz white-noise burst (70-80 dB sound pressure level) from one of two speakers, each located in front of the tangent screen and 10 deg. to the right or to the left of the fixation point.

The CPVA task consisted of up to 8 classes of trials, pseudorandomly interleaved. These were the same as classes 1-4 of the CP task, except that the stimuli could be visual or auditory. Specifically, 2 classes were memory saccade to left and right visual targets; 2 classes were memory saccade to left and right auditory targets; 2 classes had the A-B stimulus sequence, as in the CP task class 3, but with the modalities visual-then-auditory or auditory-then-visual; and 2 classes had the B-A stimulus sequence, in the visual-then-auditory or auditory-then-visual order. Four of the CPVA classes thus required a single change of plan, sometimes prompted by a visual stimulus and sometimes by an auditory one. The timing of the stimuli was slightly different in the CPVA task from the CP task. In all classes the first stimulus appeared after 750 ms of steady fixation, remained on for 750 ms and was followed by a delay (M period) lasting 750 ms. The second stimulus, if present, also appeared for 750 ms and was followed by a 750 ms delay.

Data analysis

The periods of interest in all tasks were the M periods following presentation of the stimuli. For analysis purposes we defined an M period as starting 100 ms after the offset of a stimulus and lasting until the extinction of the FP or appearance of the next stimulus, whichever came first. M periods thus lasted 300 ms in the MS and CP tasks, 1150 ms in classes 1-4 of the CPVA task, and 650 ms in classes 5-8 of the CPVA task. If a cell's response clearly started well within

one of these periods we redefined the period so that it consisted mostly of the cell's response time. For each class we computed the firing rate during M periods and compared it to the firing rate during the background period (BG) for that class. The BG period was a portion of the time preceding stimulus appearance, while the monkey fixated on the FP. It started at 300 ms from trial onset and ended at the appearance of the first stimulus (at 800 ms in the MS and CP tasks and 750 ms in the CPVA task).

We classified a neuron as excitatory if its M activity in the MS task was significantly greater than BG (two-tailed paired t test, $p<0.05$) in at least one direction for a given target eccentricity. If the neuron was *not* excitatory and had M activity significantly (statistically) smaller than BG we classified it as inhibitory.

We conducted one-tailed paired t-tests to detect significant differences ($p<0.05$) between M and B periods in the various classes of the CP and CPVA tasks.

As a quantitative measure of a neuron's change in activity from one memory period to the next we computed, for classes 3 and 5 of the CP task, the index I=(1-M2/M1), where M1 and M2 denote the net memory activity following the first and second target, respectively.

RESULTS

Database

Our database consists of 81 neurons isolated in area LIP in 3 hemispheres of two monkeys while the animals were performing the CP and CPVA tasks. The units were selected according to the following criteria: first, assignation to area LIP (as described in the Methods section of the accompanying study (previous chapter); second, presence of significant (paired t test, $p<0.05$), spatially specific, M period activity in memory saccade tasks. These neurons are a subset of a larger group of neurons isolated in area LIP of our monkeys. Because we were also conducting related experiments in parallel with those reported here (Barash et al., 1991a, b; Bracewell et al., 1991;

Mazzoni et al., 1993, 1994), not all the eligible units that we isolated were tested in the CP and CPVA paradigms. 62 neurons were excitatory and, 19 inhibitory.

We describe below the response patterns of the neurons we tested in the CP and CPVA tasks. For each task we will illustrate the response pattern of a typical neuron, followed by quantitative analysis of the responses of all the neurons in our sample.

Visually cued change of plan: examples

We tested all neurons on classes 1-4 of the CP task. Fig. 2 (a-d) shows the response pattern of an excitatory LIP neuron in these classes. This neuron had a RF in the left visual field. Fig. 2a shows that target A evoked a clear memory period activity, whereas target B did not (Fig. 2b). In Fig. 2 c we see the neuron's response in trials requiring a change of plan. Initially target A evokes a sustained response, which we suggest reflects the monkey's intention to make a saccade to A within the RF of the cell when the fixation spot is extinguished (as in Fig. 2 a). However, while he is waiting, target B appears, indicating to the monkey that he must change his plan and prepare a saccade to B outside the neuron's RF. The eye movement traces indicate that this is what he does, after the FP is extinguished. Because he maintains fixation on the FP for as long as it is present, changes in firing rate cannot be attributed to eye movements. Target B appears to "cancel" the sustained M activity evoked by target A: the clear activity present during the first M period (M1) is absent in the second (M2). Conversely, in those trials in which target B is followed by target A the neuron shows the opposite pattern of memory period activity (Fig. 2d). Target B, which is outside the receptive field, evokes no M activity (period M1). However, when target A appears briefly within the receptive field, indicating that the monkey must change his plan and program a saccade to A and not B, sustained activity is evoked (period M2).

We tested a subset of neurons (n=31) on all 8 classes of the CP task. For these neurons trials of all 8 classes were pseudorandomly

interleaved. Classes 5 and 6 were designed as control classes. The reduction in activity from the first (M1) to the second (M2) memory periods in class 3 might be due to processes unrelated to the change in saccade plan. For example, a neuron may simply adapt to repeated presentations of a visual stimulus. With classes 5 and 6 we tested whether any changes in M2 activity relative to M1 seen in change of plan trials were a nonspecific result of the fact that a target had recently appeared when the second target was presented. In these classes the target appeared twice in the same spot (at A in class 5, and at B in class 6), instructing the monkey to keep his saccade plan unchanged. In these trials the first plan formulated by the monkey was indeed the correct one. The second target therefore did not serve to change the monkey's plan. The appearance of the second target does not substantially alter the memory activity in either class of trial (Figs. 2e-f), reflecting the monkey's maintenance of the same saccade plan.

In the final two classes (classes 7 and 8) we presented three targets in sequence while the monkey maintained fixation (A then B then A, or B then A then B). In these trials the monkey was rewarded for making a saccade to the location of the last target to appear (as always); thus he presumably changed his plan twice as he was not aware how many stimuli would appear in the trial. Figs. 2g and 2h illustrate the responses of the same LIP neuron in these classes. In class 7 (Fig. 2g) we see that target A, which appears in the receptive field, evokes memory activity (during period M1) which is cancelled by the appearance of B outside the receptive field (little activity in M2) only to reappear after A is presented once more (period M3). The converse pattern of activity is observed in class 8 (Fig. 2h), in which the sequence of targets is B-A-B.

We observed the same pattern of M period activity in neurons that had inhibitory M responses. Fig. 3 shows the activity of such a neuron in the CP task. The cell's activity is inhibited by a stimulus in the upper left quadrant (target A; Fig. 3a) and is unaffected by a stimulus in the lower right quadrant (target B; Fig. 3b). In CP trials, the inhibition appears every time a saccade to target A is planned: in M1

of classes 3, 5, and 7 and M3 of class 7 (Figs. 3c, e, g). The inhibition is maintained when the plan for the up-left saccade is maintained (M2 of class 5; Fig. 3e), but is promptly cancelled every time the planned saccade is changed to a down-right one (M2 of classes 3 and 7, M3 of class 8; Figs. 3c, g, h).

Quantitative analysis

The M activity of the neurons of Figs. 2 and 3 is correlated with the monkey's saccade plan. It is significantly raised for the neuron of Fig. 2, and significantly depressed for the neurons of Fig. 3, whenever the monkey plans a saccade towards the neuron's RF. The activity of both neurons remains at background level whenever the saccade plan is for the opposite direction. The hypothesis that the M activity reflects the current saccade plan (motor plan hypothesis) predicts that this activity should change from one M period to the next within A-B, B-A, A-B-A, and B-A-B trials, and remain the same in A-A and B-B trials. Specifically, according to the motor plan hypothesis, M1 and M2 should be significantly different in classes 3, 4, 7, and 8, and not significantly different in classes 5 and 6. In classes 7 and 8, moreover, the activity in the M3 period should return to the same level as in M1. The M activity of the majority of neurons was consistent with this hypothesis in all classes of the CP task (table 1; paired t tests, $p<0.05$).

For several neurons the M3 response of class 7 was even stronger than in M1. We considered this consistent with the motor plan hypothesis. A possible explanation for this result may lie in the design of the CP task. Because no trials in this task had more than 3 target presentations, the monkeys were likely to know that the third target was certainly the goal of the required saccade. The enhancement of the M3 response over the M1 response may reflect the additional assurance that the motor plan of the M3 period would actually be executed

In order to quantitate for each neuron how much its M activity changed or stayed the same in the presence or absence of a plan change, we computed an "activity-change" index. We defined this as 1-

M2/M1, where M1 and M2 are the net responses, relative to the background activity level, during the first and second M periods of a given class, respectively. If the response turns off from one M period to the next the index approaches the value 1 (100% change of activity). If the activity remains the same the index approaches the value 0. Values greater than 1 indicate that the response reverses sign (from inhibitory to excitatory or vice versa) from M1 to M2, and values less than 0 mean that the memory response becomes even stronger (activity increases for excitatory cells and decreases for inhibitory cells) in M2 than in M1.

We computed the activity-change index for class 3, in which the plan changes from a saccade in the neuron's PD to one in the opposite direction, and for class 5, in which the plan remains unchanged as a plan for a saccade in the PD. The index values for class 3 cluster around 1 (median = 1.1; Fig. 4a), indicating that as a population the change in plan for a movement inside the RF to outside the RF cancels most of the response. Comparison of the two memory periods of class 3 revealed a significant difference in activity for 80% (65/81) of these neurons ($p<0.05$). The activity of most neurons thus reflected the change in saccade plan. In class 5, on the other hand, the index values cluster around 0 (median = 0.0; Fig. 4b), no value being greater than 0.7. Among the neurons tested, 71% (22/31) showed no significant change in activity in class 5 ($p<0.05$). The activity of most neurons tested thus remained largely unchanged in class 5, remaining consistent with the maintenance of the same oculomotor plan.

Visual-auditory change of plan

We have recently found that many units in area LIP that respond to visual stimuli in memory saccade tasks also respond when auditory stimuli are used as cues to the location of the saccade goal (Bracewell et al., 1991a, b; Mazzoni et al., 1993). Several of these neurons have a similar pattern of responses in the visual and auditory versions of a memory saccade, that is, they respond in the same phases of the task (sensory, delay, and saccade periods) and with the same spatial selectivity regardless of stimulus modality. We predicted that if the M

activity of these bimodal neurons reflects the monkey's saccade plan, then it should reflect changes of plan whether these are cued by visual or auditory stimuli.

We tested this prediction on a subset of LIP neurons (n=22) that had clear M responses in visual and auditory memory saccade trials. We recorded the activity of these neurons in the CPVA task. The activity of one of these neurons during four trial types of this task is shown in Fig. 5. This neuron had a significant M period response in rightward memory saccade trials, whether they were instructed by a visual cue (Fig. 5a) or by an auditory one (Fig. 5b). In the next two trial types the monkey had to change his plan once. In Fig. 5c the first stimulus is auditory and falls outside the neuron's RF, eliciting no M response. The second stimulus is visual and falls in the RF, eliciting clear M activity. In Fig. 5d the first stimulus, outside the RF, is visual and produces no response. The second stimulus is auditory, in the RF, and it evokes strong M activity. The neuron's activity can thus be altered, without any overt behavior, using spatial cues from two modalities.

The responses to auditory and visual stimuli of the bimodal LIP neurons often differed in strength and tuning sharpness (unpublished observations). This may have been due to the fact that we did not attempt to match the saliency of stimuli in the two modalities. These differences precluded a quantitative analysis of the responses in the CPVA task. Among the 22 neurons tested, activity in 8 clearly reflected change in plan; activity in a further 8 was in partial support of our hypothesis. Cells with strong M activity gave the clearest results in this paradigm.

DISCUSSION

We have previously shown (Gnadt and Andersen, 1988; Barash et al., 1991a, b; Bracewell et al., 1991; Mazzoni et al., 1994) that the "memory" activity exhibited by an LIP neuron reflects the intention to make a saccade in a certain direction: it is a motor memory or plan,

rather than a sensory memory. The present results clearly demonstrate that alterations in the monkey's intentions, even in the absence of overt behavior, are manifested in altered LIP activity. Altogether, our experiments establish a role for area LIP that goes beyond the perceptual components of saccade execution. Its neurons express a physiological correlate of a monkey's plan to make a particular saccade. Because their activity is not obligatorily linked to the actual execution of a saccade, it does not encode a saccadic motor command, but rather the intention to execute a saccade. This intention can be altered as new information allows the monkey to update his motor plan.

We also found that the M activity of many LIP neurons is evoked regardless of the modality of the cue to the saccade goal. Many cells in LIP are thus at least "bimodal", encoding plans for saccades to behaviorally relevant targets regardless of the modality through which these targets are localized.

Controls

The monkeys had to maintain stringent fixation (within 1 deg. of the fixation spot) for as long as the fixation spot was present. Thus the alterations in activity during the M period(s) of the trials could not have been due to eye movements. Moreover, trials were pseudorandomly interleaved; thus the monkey did not know how many (nor which) targets would be presented. In order to perform nearly flawlessly (as they did), they presumably had to plan to saccade to a target when it was presented, and change their plan if necessary later during the delay period.

Finally, classes in which the same target was presented twice (A-A and B-B trials) ensured that the monkeys could not use the appearance of a second target as a non-spatial cue to change their plans. They also tested controlled for adaptation as a possible cause of response reduction from one M period to the next.

Other studies of changes of motor plan

Wise and colleagues have performed extensive studies of "motor set" (in our terminology, motor set is "motor intention") in the premotor cortex (PMC; for review, see Wise, 1985). Their most important finding is that, for most PMC neurons, the delay period activity is related to the direction of the forthcoming arm movement (the motor set), and not to the visuospatial cue ("instructional stimulus") indicating which response is required (Wise et al., 1983; Weinrich et al., 1984).

In a study in which the instructional stimulus was changed during the delay period, directionally-specific motor set units showed concomitant changes in activity (Wise and Mauritz, 1985). These results are similar to those of our "change of plan" study. One difference between their study and ours was that we did not present our targets (instructional stimuli) for the whole duration of the delay period, as did Wise and Mauritz. In theory, the changes in sustained activity observed by Wise and Mauritz (1985) might have been due to the changes in the continually present instruction stimulus. However, Wise and Mauritz (1985) also demonstrated that delay period activity did not depend on the continual presence of the instructional stimulus, which suggests that the continual presence of the instructional stimuli in their change of motor set experiment did *not* account for the change in activity with change in motor plan which they observed.

Relationship between LIP activity and saccades

Our results suggest that LIP is involved in the planning of saccades, and rather indirectly in their production. Activity in LIP is not necessarily followed by a saccade. LIP activity may vary in the absence of overt eye movement behavior. It is possible that LIP projections to superior colliculus (and perhaps to the frontal eye fields) may raise the level of excitation there such that a "trigger" signal can more easily evoke a saccade. In this regard it is interesting that LIP cells often show a saccade-related burst of spikes with a frequency substantially higher than that during the sustained, elevated M

activity (e.g., see Fig. 2a). This burst may serve as part of the trigger signal suggested above. The initial source of this saccade response is unknown and would require further experiments to determine the first locations in the brain to generate it.

The activity in other high order motor areas typically also has a non-obligate relation to movement (reviewed in Georgopoulos, 1991). This is true even of the corticomotoneuronal cells of the primary motor cortex cells, the "upper motor neurons" of many a neurology textbook. Evarts (1981;, 1986), Cheney and Fetz (1980) and Lemon (1988) have shown that the relationship between their firing and muscle activity is conditional and complex. In addition, some neurons in the motor cortex fire during an instructed delay period of a delayed response task (e.g., Evarts and Tanji, 1974). It is thus perhaps better to think of high order areas in terms of motor planning, and to "consign" the details of execution to lower regions such as the brainstem (for eye movements, see Wurtz and Goldberg, 1989 for review) and the spinal cord (for limb movements, see Alsternack et al., 1981; Georgopoulos and Grillner, 1989).

TABLE 1

Number (and percentages) of neurons whose M activity in the CP task is consistent/not consistent with the motor plan hypothesis.

	Consistent with motor plan	Not consistent with motor plan	Total
Class 3 (M2 vs. M1)	65	16	81
	(80%)	(20%)	(100%)
Class 4 (M2 vs. M1)	62	19	81
	(77%)	(23%)	(100%)
Class 5 (M2 vs. M1)	22	9	31
	(71%)	(29%)	(100%)
Class 6 (M2 vs. M1)	24	7	31
	(77%)	(23%)	(100%)
Class 7 (M2 vs. M1)	25	6	31
	(81%)	(19%)	(100%)
Class 7 (M3 vs. M1)	28	3	31
	(90%)	(10%)	(100%)
Class 8 (M2 vs. M1)	26	5	31
	(84%)	(16%)	(100%)
Class 8 (M3 vs. M1)	27	4	31
	(87%)	(13%)	(100%)

Excitatory M activity was consistent with the hypothesis if it decreased from M1 to M2 in classes 3 and 7, increased from M1 to M2 in classes 4 and 8, remained unchanged from M1 to M2 in classes 5 and 6, and returned at least to the same level from M1 to M3 in classes 7 and 8 ($p<0.05$).

FIGURE LEGENDS

Figure 1:

Spatial and timing paradigms of the CP task. The panel at the top shows the spatial arrangement of the stimuli (dots) relative to the fixation point (FP, cross) and to the neuron's receptive field (RF, dotted semicircle). The stimuli appeared either in the RF at location A or outside the RF at location B. For each class we show the appearance of a stimulus as a thick horizontal line. In all classes the FP appears first and remains on as one, two, or three stimuli appear in sequence. M, M1, M2, and M3 refer to the "memory," or delay, periods that follow each stimulus appearance.

Figure 2:

Activity of an excitatory LIP neuron in the CP task. The abscissa in each panel represents time (100 ms/division). Each panel contains, from top to bottom: rasters of tick marks representing the occurrences of action potentials, each row corresponding to one trial; a time histogram (bin width = 50 ms) of the neuron's average rate of action potential firing over all trials (25 Hz/division); and a trace of the monkey's vertical eye position (30 deg./division). Onset and offset times of stimuli during the trials are indicated both by the thin vertical lines within each panel and by the thick horizontal lines below each panel. Abbreviations are as in Fig. 1a. *a-b)* Simple memory saccade towards the RF (class 1) or away from it (class 2); *c-d)*, single change of plan (A then B in class 3; B then A in class 4); *e-f)*, no change of plan (controls) (A then A in class 5; B then B in class 6); *g-h)*, double change of plan (A-B-A in class 7; B-A-B in class 8). Trials with 1, 2, and 3 targets were pseudorandomly interleaved so that the monkey could not predict the target sequence or the required saccade in advance.

Figure 3:

Activity of an inhibitory LIP neuron in the CP task. The data is presented in the same format as in Fig. 2, using the same abbreviations. The histogram's bin width is 100 ms and its vertical scale is 20 Hz/div.

Figure 4:

Index of change of neural activity between the first and second memory periods of the CP task. The index is computed as 1-M2/M1, where M1 and M2 are the average net firing rates, relative to the background rate, in the first and second memory periods, respectively. Excitatory as well as inhibitory cells are included. *a)*, Index values in class 3 for all neurons in our database (n=81. *b)*, Index values for class 5 for the neurons tested in classes 5-8 (n=31).

Figure 5:

Activity of an excitatory LIP neuron in the CPVA task. The data is presented in the same format as in Fig. 2, using the same abbreviations. The histogram's bin width is 100 ms and its vertical scale is 20 Hz/div. In *a)* and *b)* a visual and an auditory stimulus, respectively, are presented in the neuron's RF. In *c)* the auditory stimulus appears outside the RF and the visual one inside the RF, while in *d)* the visual stimulus is outside the RF and the auditory stimulus inside.

CP Task
A
B
RF.
FP
500 ms
Class 1
A
M
FP
Class 2
FP
B
M
Class 3
A
M1
M2
B
FP
Class 4
A
M2
M1
B
FP
Class 5
A
M1
A
M2
FP
Class 6
FP
M1
M2
B
Class 7
A
M1
A
M3
M2
B
FP
Class 8
A
M2
M1
M3
B
FP

Figure 1

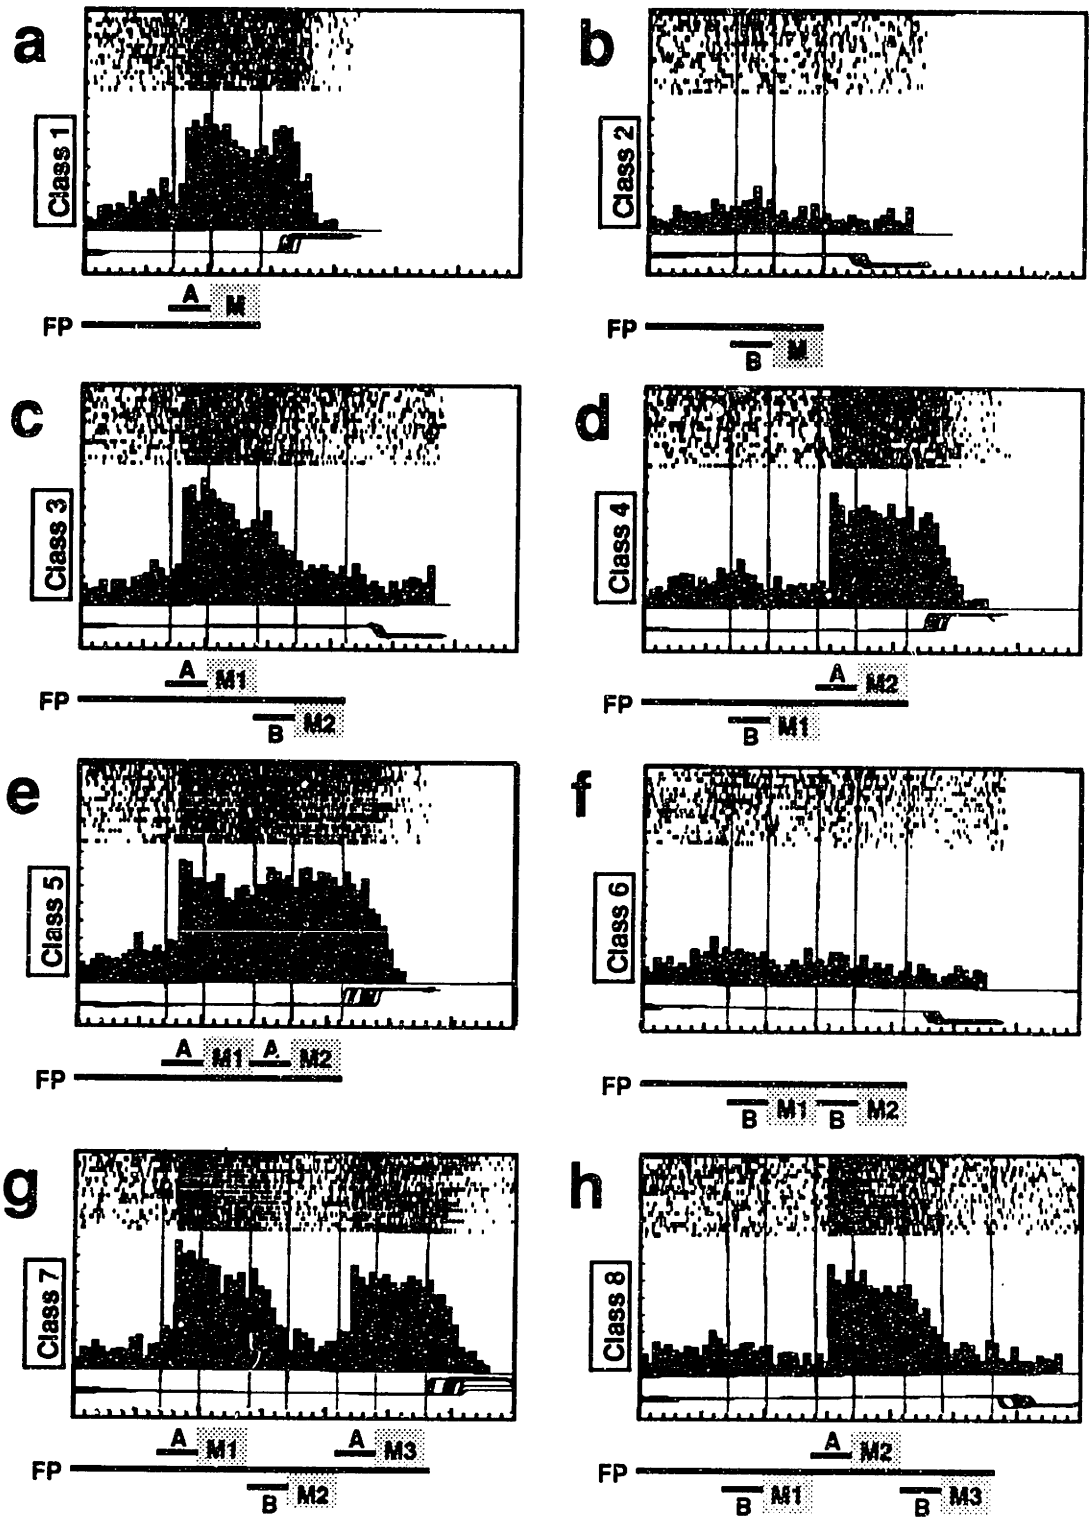
a
Class 1
A
M
FP
b
Class 2
FP
B
M
c
Class 3
A
M1
FP
B
M2
d
Class 4
A
M2
FP
B
M1
e
Class 5
A
M1
A
M2
FP
f
Class 6
FP
B
M1
B
M2
g
Class 7
A
M1
A
M3
FP
B
M2
h
Class 8
A
M2
FP
B
M1
B
M3

Figure 2

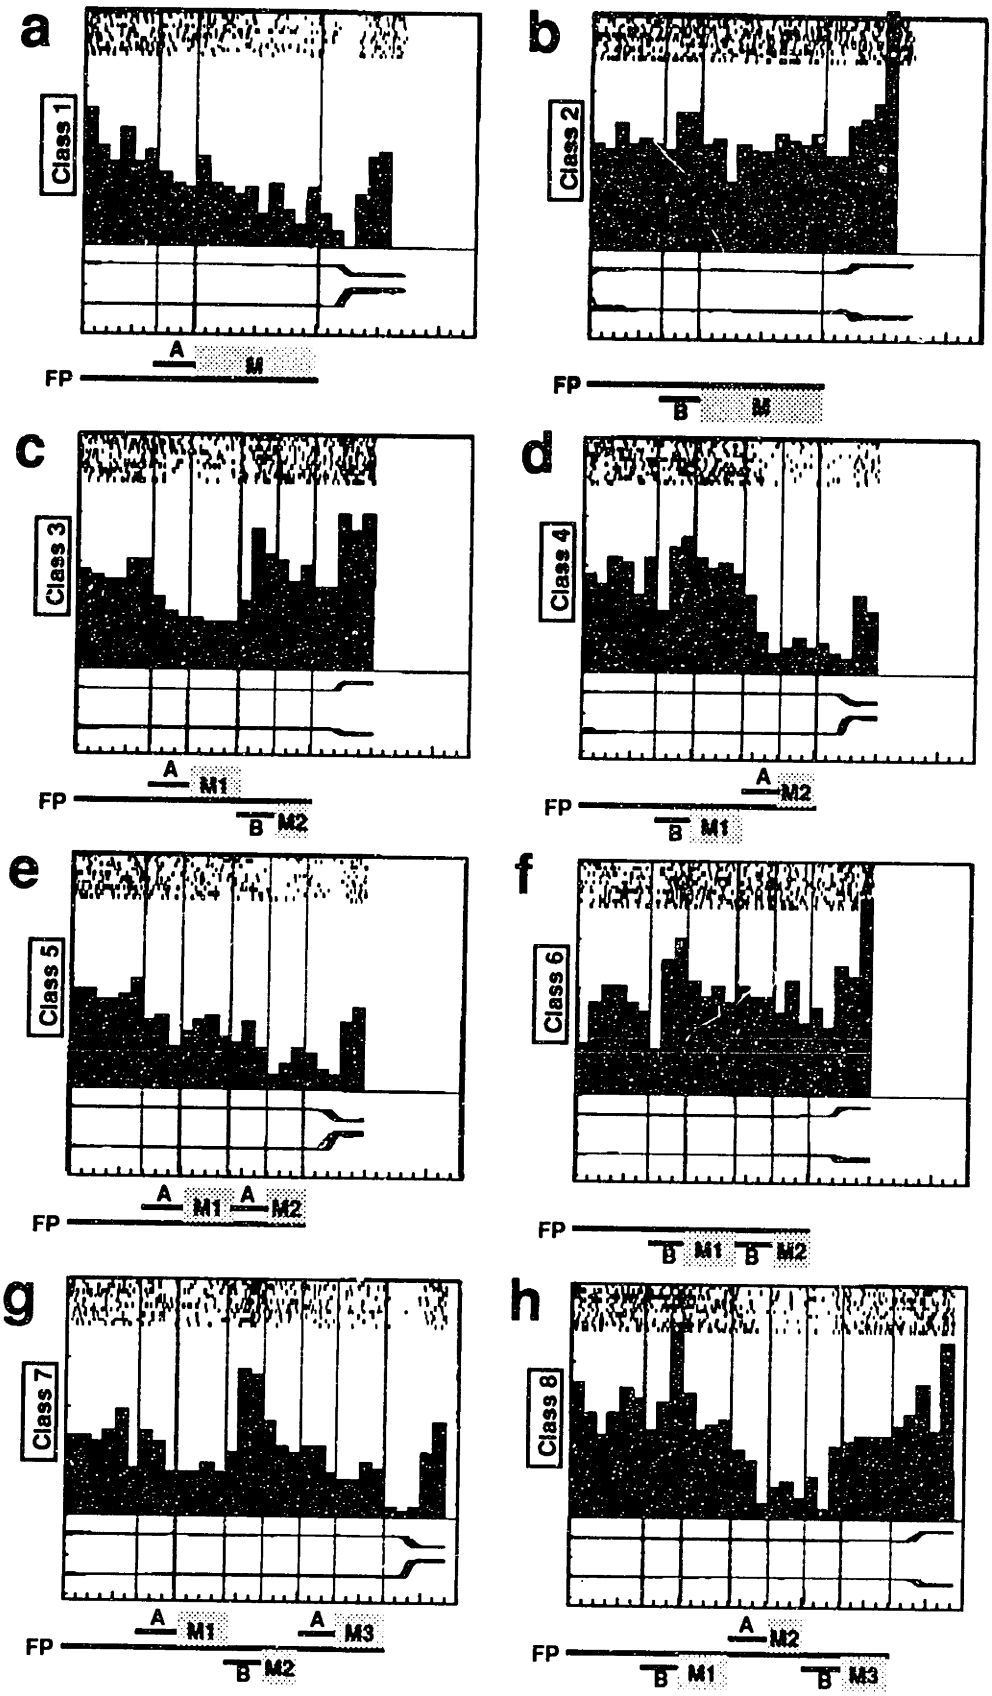
a
Class 1
A
M
FP
b
Class 2
FP
B
M
c
Class 3
A
M1
FP
B
M2
d
Class 4
A
M2
FP
B
M1
e
Class 5
A
M1
A
M2
FP
f
Class 6
FP
B
M1
B
M2
g
Class 7
A
M1
A
M3
FP
B
M2
h
Class 8
A
M2
FP
B
M1
B
M3

Figure 3

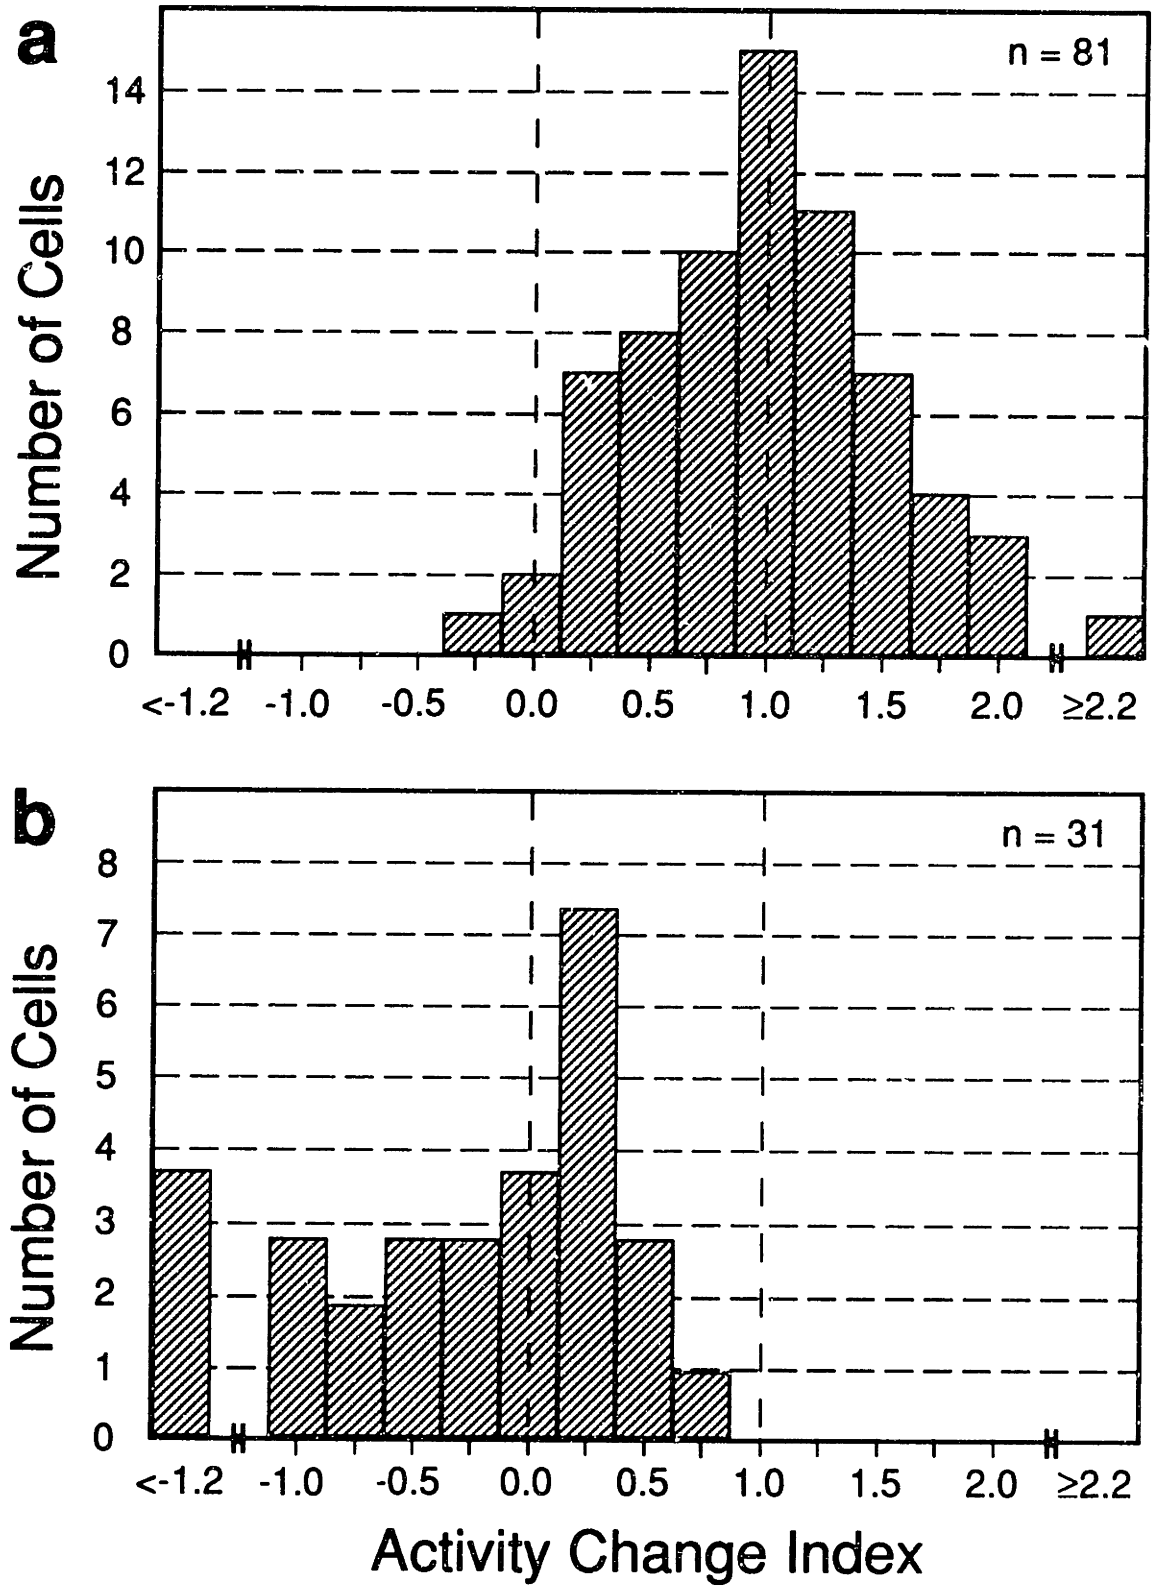
a
n = 81
Number of Cells
14
12
10
8
6
4
2
0
<-1.2
-1.0
-0.5
0.0
0.5
1.0
1.5
2.0
≥2.2
b
n = 31
Number of Cells
8
7
6
5
4
3
2
1
0
<-1.2
-1.0
-0.5
0.0
0.5
1.0
1.5
2.0
≥2.2
Activity Change Index

Figure 4

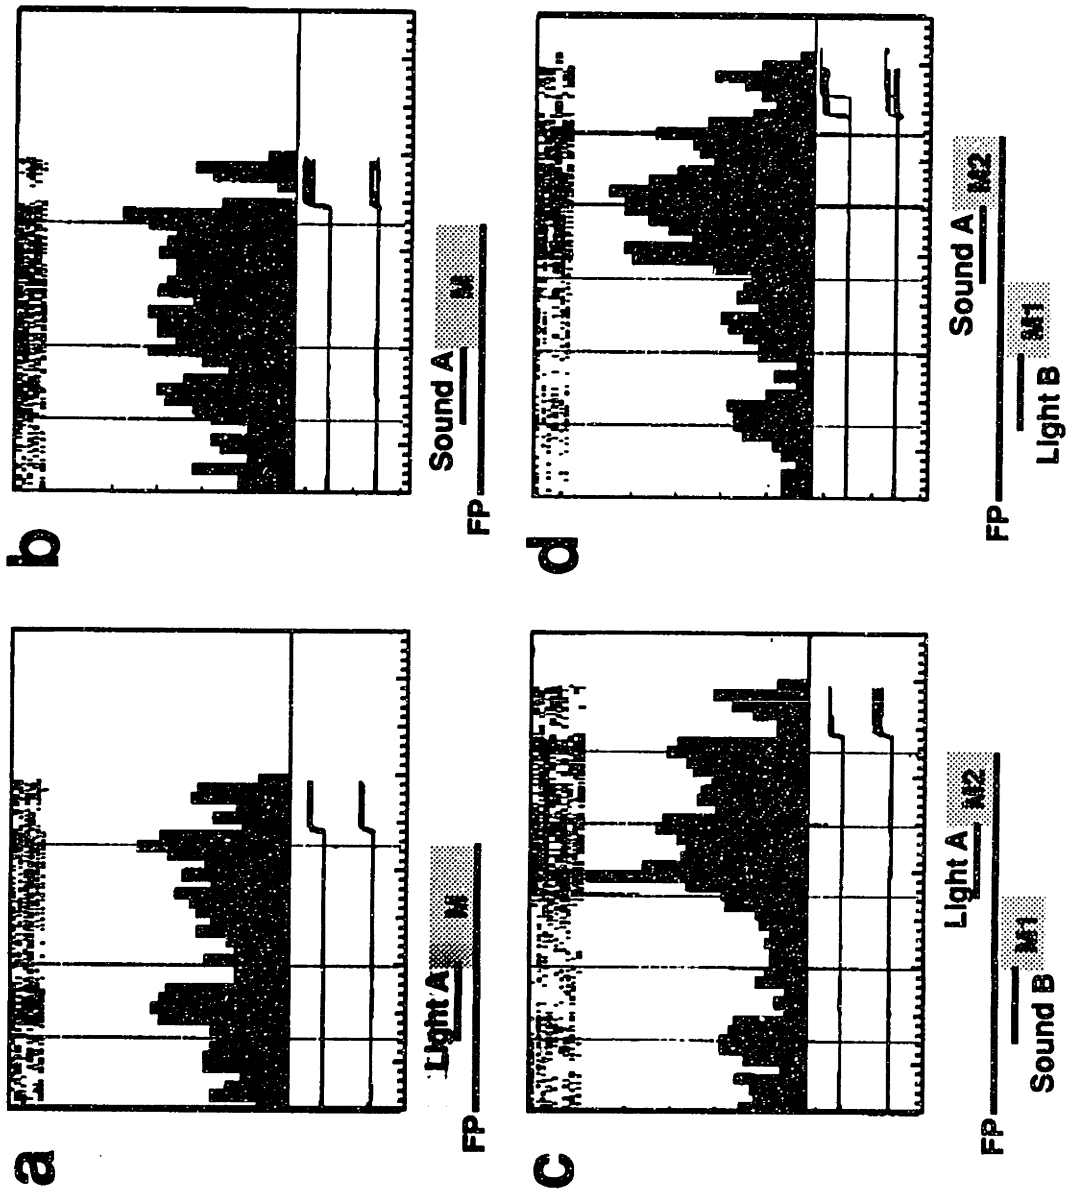
a
Light A M
FP
b
Sound A M
FP
c
Light A M2
Sound B M1
FP
d
Sound A M2
Light B M1
FP

Figure 5

Chapter 9

Effects of reversible lesions of posterior parietal cortex on saccadic eye movements

SUMMARY

The responses of neurons in the posterior parietal cortex (PPC) described in previous chapters suggest that this brain structure (and especially area LIP) plays an important role in the programming of saccadic eye movements. I have thus inactivated portions of a monkey's PPC to see what effects this intervention would have on saccadic eye movements. In order to reproduce this experiment more than once in the same animal I made reversible lesions by injecting the neuronal activity inhibitor muscimol (an agonist of γ-amino-butyric acid, GABA). To test whether the difficulty of the type of saccade made any difference on any effect I might observe, I tested the monkey's ability to perform visually guided saccades as well as saccades to the remembered locations of visual targets (memory saccades).

Visual saccades were only mildly affected. There was a small rightward shift of the saccadic endpoint for the contralateral and vertical directions (3 deg.; $p<0.05$), and an increase in latency for all contralateral and vertical directions and in one ipsilateral direction ($p<0.05$; max. increase=200%). The relationship between saccade velocity and amplitude remained unchanged.

The metrics of memory saccades were grossly altered. There was a significant rightward shift of saccadic endpoint for saccades in all directions ($p<0.05$). This shift was larger for contralateral and vertical saccades (average=12 deg.) and smaller for ipsilateral saccades (average 5 deg.). Latency was significantly increased in only two of the contralateral directions. The velocity of contralateral and upward memory saccades was reduced to an extent indicating a shift toward lower velocities of the velocity-amplitude relationship.

These results suggest that the PPC plays is essential to the correct programming of memory saccades but not of visually guided saccades. The PPC appears thus to play a unique role in programming orienting movements that are more complex than those under direct visual guidance.

INTRODUCTION

The theoretical and experimental studies described in the previous chapters strongly implicate area LIP in saccadic eye movements. Such studies, however, can only establish correlations between neural activity and behavior. The modelling studies show that area LIP can transform the coordinates of sensory stimuli into different reference frame, and the recording studies show that area LIP carries a code of the monkey's saccade plan. These results alone, however, do not establish that area LIP's activity is necessary to make saccades. Lesion studies can address this question.

Studies of neurological patients were among the first to implicate the posterior parietal lobe in the control of saccadic eye movements. [Although Ferrier (1876) had elicited eye movements by electrically stimulating the PPC of anesthetized monkeys, he had taken this result as evidence that the PPC was involved in visual processing, and not in eye movement control.] Besides producing the classic syndrome of unilateral neglect (Critchley, 1953; Bisiach and Luzzatti, 1978; De Renzi, 1982; Heilman et al., 1987), unilateral lesions of the posterior parietal lobe in humans also cause mild deficits of saccadic eye movements (Holmes, 1938; Critchley, 1953; De Renzi, 1982). Bilateral lesions, on the other hand, produce great difficulty in making saccades (Balint's syndrome; Balint, 1909; Hécaen and De Ajuriaguerra, 1954; reviewed in Husain and Stein, 1988). In the latter condition, patients display remarkable difficulty in intentionally shifting their gaze from one direction to another, while reflexive and spontaneous eye movements are relatively preserved—a symptom which Balint termed "psychic paralysis of gaze."

Previous studies of PPC lesions (unilateral as well as bilateral) in monkeys reported some impairment in the performance of visually guided saccades (Lynch, 1980; Lynch and McLaren, 1989), notably an increase in latency and some dysmetria for certain saccade directions. These deficits are generally milder, however, than one might predict based on anatomical and physiological evidence. The output of area

LIP constitutes a major source of input for both the superior colliculus (SC) and the frontal eye fields (FEF). These structures are well-established saccade-control centers (reviewed in Schiller, 1984; Sparks et al., 1989; Goldberg and Bruce, 1989), and the removal of *both* of them virtually abolishes saccadic eye movements in monkeys (Schiller et al., 1980). Given that the signals that LIP neurons send to the SC and FEF are intimately related to saccades (reviewed in Andersen, 1989; see also Chapters 5-8), one might expect a major disruption of saccades to result from removal of the PPC.

Two reasons why the deficits obtained in monkeys are milder than one might expect from the type of processing the PPC appears to perform, could be technical ones. Permanent lesions usually involve a period of several hours during which the monkey cannot be tested on a given task. This period may be sufficient for enough recovery to occur to mask immediate deficits. The task itself, moreover, may not be one that taxes the truly relevant function of the area in question.

I addressed these possibilities by making reversible lesions of PPC via injection of the GABA agonist muscimol and by testing the monkey on memory saccades as well as visually guided saccades. The technique of injecting an anesthetic virtually eliminates the confounding influence of recovery—at least recovery that requires a few hours to take place. Using memory saccades, besides visual saccades, as probes of gaze-shifting abilities provides a more complex task in which deficits might be more obvious than in visually guided saccades. Our preliminary results include a specific disruption of the metrics of memory saccades, with much smaller changes in visually guided saccades. These results suggest that the PPC may play a special role in guiding orienting movements that are not under direct visual control.

METHODS

Experimental set-up, animal care, data collection, and eye movement tasks

These components of our experiment were identical to those described in Chapter 7. We used one of the same monkeys (*Macaca mulatta*) as in the experiments described in Chapter 7. The behavioral tasks used in this experiment were the *visual saccade* and *memory saccade* tasks. The parameters of these tasks were as described in Chapter 7 with the following exceptions: 1) the window of acceptance of saccade endpoint was enlarged to a square of 10 deg. side around the target's position for the visual saccade task and a square of 40 deg. side for the memory saccade task, so that even grossly dysmetric saccades could be stored for analysis; 2) the allowed latency (i.e. the period after fixation point offset until which a trial was aborted if the monkey did not make a saccade) was set at 2000 ms to allow for possibly increased saccade onset latencies.

Because some of the lesions produced nystagmus (see Results), we also recorded the monkey's eye position in a *fixation task*. In this task the monkey had to maintain fixation on a light spot that appeared at one of 9 possible positions on the screen. These positions were arranged as a square grid of 20 deg. spacing, centered on the straight ahead direction. The time of required fixation and the size of the window of accepted eye positions around the fixation spot were varied—between 300 and 1000 ms, and between 10 and 40 deg., respectively—in order to obtain records of eye position traces in spite of nystagums and deficit in the ability to fixate in certain directions (see Results).

Other behavioral tests

Besides recording eye movements in the saccade paradigms described above, I administered the following set of qualitative behavioral tests. The monkey sat in a primate chair with his head fixed during these tests, which were conducted in a well-lit room. I

recorded the monkey's behavior on videotape, and then studied the videotape for gross for gross behavioral changes.

I tested for neglect by presenting positive and aversive stimuli in the left and right hemifields around the midline of the head and observing the monkey's reaction. Positive stimuli were raisins, which the monkey unhesitatingly reached for, grasped, and ate before the series of injection experiments. The aversive stimulus was the presentation, in the monkey's field of view, of a 1 cc syringe like the ones used to anesthetize the monkey before any minor painful procedure. Before any lesions were performed, the monkey always became visibly agitated at the sight of this object anywhere in his visual field. He shifted in his chair as if to withdraw from the syringe, and waved either of his arms as if to push the syringe away.

I could not use the magnetic record the monkey's eye position in these tests. I evaluated the monkey's oculomotor abilities by recording on videotape a close-up view of the monkey's eyes while I manually moved a raisin in various regions of the monkey's visual field. I tested the monkey's ability to look in different directions by quickly bringing a raisin into view, from outside the monkey's visual field, to one of 8 peripheral position in the visual field (up, up-left, left, down-left, down, down-right, right, up-right). I tested the ability to make smooth pursuit eye movements by moving a raisin horizontally and vertically across the monkey's visual field at approximately 25 deg./s.

Solutions

The neuronal inactivator used in this study was a solution of the GABA agonist muscimol (Sigma, St. Louis, Missouri) in saline (8 mg/mL). Injections of saline alone were used as controls.

Injection method

Solutions were injected manually via a 1 μL syringe (Hamilton, Reno, Nevada) mounted on a microdrive (Narishige, Tokyo, Japan) specially adapted to manipulate a Hamilton syringe. At the beginning of each session the syringe's needle was inserted in the microdrive's

guide tube and the latter was pushed through the dura mater. The needle was then advanced to a chosen depth using the microdrive's motor. Injection sites were chosen based on previous recording of the activity of LIP neurons, performed during the studies described in Chapter 7.

RESULTS

I made several injections of muscimol as well as saline in one monkey. All injection sites were chosen based on recording studies that had taken place during several months *prior* to the lesion study described here. I did not record neuronal activity at the injection site before or after the injection. I have converging indirect evidence (discussed below), however, that I inactivated at least portions of the PPC. At this point I cannot know which areas within (and outside) the PPC were inactivated, nor which were responsible, when inactivated, for the various effects I observed. My attribution of these effects to inactivation of the PPC is tentative, and the following results are to be thus considered preliminary. I can exclude inactivation of axons or other fibers of passage as sources of the effects observed because muscimol is an agonist of GABA receptors (Andrews and Johnson, 1979), which are not found on these structures.

Injections

Injection of muscimol produced clear effects on saccades at three sites in the right PPC. For each of these sites, the coordinates of the penetration in the monkey's chamber corresponded to sites where LIP activity had been recorded during the experiments described in Chapter 7. At one of these sites (to which I will refer as site G) I obtained similar effects twice on separate days. I also injected an equal amount of saline at this site on a separate day as a control, and obtained no clear effect. I report below the effects on saccades obtained by injection of 1 μL of muscimol (8 μg/μl) at site G.

Effects on visually guided saccades

Fig. 1 shows the vectors of saccades to visual targets presented at a 15 deg. radius along 8 different directions. The clearest effect of muscimol injection is a rightward shift of the endpoints of upward and downward saccades. This effect lasted for several hours but disappeared by the next day (Fig. 2).

Quantitative comparisons revealed a significant rightward shift of saccade endpoint in each of the contralateral and vertical saccades, averaging about 3 deg. ($p<0.05$) (Fig. 2). The vertical components of the saccades decreased slightly, by about 1 deg., for the directions 90, 45, and 180.

Saccade latency was significantly increased ($p<0.05$) for all contralateral and upward directions, as well as for the 45 deg. direction (Fig. 3a), reaching a two-fold increase for the 180 deg. direction.

Peak saccadic velocity decreased ($p<0.05$) for the contralateral directions and the upward direction (Fig. 3a). This decrease is expected, however, given the reduction in amplitude produced by the reightward shift of enpoints of contralateral saccades. When plotted against saccade amplitude (Fig. 4), the peak velocities of saccades following injection appeared to fall on the same main sequence curve as the one outlined by the peak velocities before the injection.

Effects on memory saccades

The main effect of injection on memory saccades was also a rightward shift of saccade endpoint. This was much more pronounced, however, than for visual saccades, and occurred for contralateral as well as ipsilateral saccades (Fig. 5). This effect also virtually disappeared the next day (Fig. 6).

Quantitative comparisons, for each direction, of the horizontal and vertical components of the memory saccades revealed a significant rightward shift of the horizontal component for every saccade direction (Fig. 7a), which amounted to about 4 deg. for the ipsilateral and vertical directions and reached 12 deg. for directions 180 and 225. The

vertical component, on the other hand, was virtually unaffected, being significantly decreased (by 3 deg.) only for vertical saccades (Fig. 7b).

There was an increase in latency for the contralateral directions (significant for directions 135 and 225), and a decrease for some of the ipsilateral directions (significant for rightward saccades) (Fig. 8a).

Peak velocity decreased for almost all saccade directions (significantly in direction 0-225 deg.; Fig. 8b). This was expected for the contralateral directions given that, for these directions, as rightward shift of the endpoint produced a reduced saccade amplitude (and thus a reduction in peak velocity according to the main sequence curve). The same rightward shift of enpoints, however, also resulted in an increased amplitude for ipsilateral saccades. This increase was not accompanied by any significant increase of peak velocity that would be expected from the main sequence curve, suggesting that for memory saccades muscimol injection also affected the relationship between amplitude and velocity. This shift of the main sequence curve is relatively clear in Fig. 10.

Recovery

The effects of muscimol injection on saccadic eye movements lasted several hours (from 5 to 12). They invariably recovered by the next day (e.g. see Figs. 2, 6). ANOVA revealed no differences, except for very small ones in a few contralateral directions, in horizontal or vertical saccade components, latency, or velocity between the pre-lesion and recovery conditions ($p<0.05$).

Effect of darkness

Performance on visually guided saccades was tested in total darkness so that the monkey could use as little information as possible besides the spatial cues provided by the fixation point and the target. Most of the memory saccade runs, on the other hand, were performed with the room lights turned on. Turning the lights off strongly exacerbated the memory saccade deficits to a point that made it difficult to collect sufficient data for analysis. Qualitatively, memory

saccades in total darkness had even more right-shifted and scattered endpoints, and increased latencies compared to the condition with the room lights on.

Controls

As a control I injected 1 μL of saline at site G in a separate session. This produced no clear changes in the performance of visual or memory saccades.

Effects of single injections at other sites

Single injections of muscimol at a site (in the right hemisphere) c. 2.0 mm away from site G, approximately in the anterior direction, produced a continuous nystagmus (Figs. 11-12) combined with a difficulty in moving the eyes into the lower-left (contralesional) quadrant (Fig. 12). The slow phase was approximately in a 45 deg. direction toward the up-left visual field. The speed of the slow phase was approximately constant within each cycle but depended on the task condition, the eye position, and the room illumination. It was always higher in the following conditions: 1) when the monkey was performing a fixation task, compared to the periods between sets of trials, when no particular behavior was required (the "no-task" condition; the monkey often became drowsy in these periods); 3) as the target light spot for fixation appeared at more downward and contralesional positions than upward and ipsilesional positions; 2) when the room lights were on, compared to darkness (in both the fixation task and the no-task conditions).

The nystagmus started within 20 min. after the muscimol injection and continued for several hours (6-13). The monkey seemed unable to stop it throughout this period. He attempted to maintain fixation on a target by making repeated saccades (perhaps better described as fast phases in this case) back to the fixation target to compensate for the drift away from it during each slow phase. Due to the dependence of slow-phase velocity on eye position, the monkey was able to maintain gaze within a relatively small range of eye positions around the target

when the latter was in the upper-ipsilateral visual field. For target positions that were more contralateral and downward the slow-phase velocity increased and the eyes swung through a larger and larger range within each cycle as the monkey attempted to maintain fixation (Fig. 12).

Simultaneously with the nystagmus, a difficulty in directing gaze towards the lower-contralesional visual field also became apparent (Fig. 12). The average position of the monkey's eyes was shifted in the right and upwards direction as the monkey attempted to fixate targets in the lower-left quadrant. This deficit was worse for the one target located at (-20, -20) deg. from the center position, but was also present (though less pronounced) for the targets in the leftward direction (at -20, 0) and in the downward direction (at 0, -20; Fig. 12). This deficit was also exacerbated by darkness.

Effects of multiple injections

In order to explore the effects of inactivating large portions of the PPC I made, in three separate sessions, muscimol injections at multiple sites and depths around the intraparietal sulcus (whose location was roughly defined by previous recording sessions). In two of these sessions I observed several symptoms of the hemineglect syndrome.

When presented with a raisin in the left (contralesional) hemifield, the monkey reached for it only after a few seconds' delay. He usually moved his right arm with relative ease up to the midline, after which he proceeded clumsily, groping for the raisin until he grasped it. Such deficits were not apparent when he reached for a rasin in the left (ipsilesional) field using his right arm.

The monkey used his left arm reluctantly and more clumsily than his right arm. When forced to use his left arm (by manually holding his right arm next to his trunk) to grasp a raisin, he did so slowly and clumsily in both hemifields, but more easily in the right hemifield than in the left one.

When an aversive stimulus (a syringe) was presented in the right hemifield the monkey immediately became agitated and started moving his right arm repeatedly in an attempt to push the syringe away from view. When I presented the same stimulus in the left hemifield the monkey completely ignored it. As I slowly moved the stimulus toward the right hemifield, the monkey continued to ignore it until it just crossed his midsagittal plane, at which point he seemed to suddenly notice it and became visibly agitated.

These symptoms appeared not to be limited to visual stimuli. When I called the monkey's name from a few centimeters away from his ear, while standing behind him out of view, the monkey was hardly startled when called from the left side compared to the right side. Similarly, while touching his head near the ear on the right made the monkey push my hand away with his arm, touching his head on the left side rarely elicited any reaction.

I examined the reference frame in which these neglect symptoms were based by rotating the monkey's head relative to the his shoulder by approximately 15 deg. to the right and to the left of the trunk's midline. The vertical plane dividing the space around the monkey into a region in which stimuli elicited many fewer reactions than in the other moved with the head; i.e., the deficits were for a hemifield defined in head-centered coordinates.

DISCUSSION

Inactivation of a portion of the PPC produced a marked impairment in the metrics of memory saccades, with much smaller effects on visually guided saccades. The most notable effect was a rightward shift of the endpoints of memory saccades, which was largest for contralateral saccades but was also clear for ipsilateral saccades.

These results are preliminary because I have performed this experiment on one monkey alone, and because the inactivated region was not confirmed by electrophysiological recording and has not been identified anatomically. Our target area for inactivation was area LIP,

but our injections could have inactivated it as well as several other areas around the intraparietal sulcus. These results are remarkable, however, in their consistency with the neural activity observed in area LIP. In the memory saccade task, LIP neurons have prominent memory period and presaccadic responses (Gnadt et al., 1988; Andersen et al., 1990b; Barash et al., 1991a, b; see also the previous Chapters). These responses are spatially tuned and encode (during the memory period) the next planned saccade. Moreover, the coarse spatial tuning of LIP neurons' activity (Barash et al., 1991b) suggests that this plan is encoded by the simultaneous activity of a population of neurons with a variety of preferred spatial directions. Finally, the spatial preferences of LIP neurons in each hemisphere are biased toward the contralateral directions (Barash et al., 1991b). Altogether, these features are consistent with a model in which area LIP maintains the plan for the next saccade in the form of the activity of many neurons in both hemispheres, with the two hemispheres acting in a push-pull fashion due to the bias of their neurons' preferred directions. This model would predict that inactivating area LIP in one hemisphere would modify memory saccades of *all* directions by shifting their endpoints in the ipsilesional direction.

Deficits specific to memory saccades are thus entirely consistent with inactivation of area LIP. Another area that has memory-period and saccade-related responses in the PPC is area 7a, which is adjacent to area LIP (Andersen et al., 1985a, 1990a). The memory activity of area 7a neurons is—at least over a population of neurons—weaker than in area LIP, and most of the saccade-related responses occur later than those of LIP neurons and after the saccade (Barash et al., 1991a). Area 7a, moreover, is much less strongly connected to eye-movement-related structures such as the frontal eye fields (FEF) and the superior colliculus (SC) than area LIP (Lynch et al., 1985; Andersen et al., 1985a, 1990a; Blatt et al., 1990). These features suggest a less direct role of area 7a than area LIP in saccadic control, making LIP a better candidate for the region whose inactivation produced the memory saccade deficits observed.

More than one cortical area may have been inactivated by each muscimol injection. It is unlikely, however, that the saccade deficits observed were due to inactivation of very large portions of PPC. No such deficits were obtained by injecting muscimol at a site that was 2.0 mm removed (anterolaterally in the chamber) from a site where inactivation produced memory saccade deficits.

Comparison with related studies

Previous studies of the effects of parietal lobe lesions on eye movements in monkeys have focused on visually guided saccades and optokinetic nystagmus (OKN). Lynch and coworkers have reported, following unilateral lesions of PPC in monkeys, a slowing of the ipsilesionally directed slow phase of OKN (Lynch and McLaren, 1983), and mild hypometria and increased latencies of contralesional visual saccades (Lynch and McLaren, 1989). Interestingly, this mild impairment of visual saccades worsens considerably when the ipsilateral FEF is also removed (Lynch 1992). The PPC seems thus to provide essential signals to the parallel pathway for saccade control going through the SC.

We also observed slight hypometria and an increase in latency of contralateral visual saccades following PPC inactivation. While the hypometria was slight, as reported by Lynch and McLaren (1989), however, the latency increase was much larger, reaching a twofold increase for the horizontal contralateral saccades. This difference may be due to our observation of immediate changes in saccade parameters. An analogous difference of results was obtained when Hikosaka and Wurtz (1985) injected muscimol into the SC. They observed much bigger changes in saccade parameters than had been observed following permanent lesions (Schiller et al., 1980; Albano and Wurtz, 1982).

The more striking effects in our study, however, were the changes in the metrics and dynamics specific to memory saccades. These support recent data obtained from studies of patients with focal cerebral lesions. These studies suggest that the posterior parietal lobe

of humans participates in the control of saccades that require more complex processing than saccades under direct visual guidance. Pierrot-Deseilligny et al. (1991) demonstrated increased latency and decreased accuracy for contralateral memory saccades in parietal patients, while Duhamel et al. (1992b) and Heide et al. (1993) have found that parietal patients have difficulty making a second saccade when it must be programmed using visual information available only before the first saccade.

Specific deficits in the memory saccade task have been observed following lesions of other structures in monkeys. The superior colliculus (SC) and the frontal eye fields (FEF) to have parallel access to the saccadic generators in the brainstem: permanent lesions of either alone produces definite but minor saccade deficits (Schiller et al., 1980; Albano and Wurtz, 1982) from which monkeys recover, whereas ablation of both structures produces an almost complete gaze paralysis from which there is little recovery (Schiller et al., 1980). The two structures do not seem to play equivalent roles in saccade programming, however: unilateral FEF ablation impairs a monkey's ability to learn to make contralateral memory saccades (Deng et al., 1986). These saccades, moreover, have irregular velocity profiles and show a depressed peak velocity-amplitude relationship. These deficits suggest an essential role of this structure in initiating and executing memory saccades. This type of saccade could in this case be representative of movements that must be triggered internally, or cognitively. A memory saccade deficit would then be consistent with the traditional view of the frontal lobe as an essential structure for initiating movements based on an internal representation or plan. It would also be consistent with the strong inhibitory projection from the FEF to the substantia nigra, a structure which projects excitatory synapses to the omnipause neurons in the brainstem and thus inhibits saccade execution (reviewed in Goldberg and Segraves, 1989, and Hikosaka and Wurtz, 1989).

Memory saccades could also be impaired if the animal simply forgot where the target had appeared. This type of deficit occurs after lesions

of the dorsolateral prefrontal cortex (Sawaguchi and Goldman-Rakic, 1991; Funahashi et al., 1993). In this study the deficits seemed purely mnemonic and not related to visual perception or saccade execution.

It cannot be distinguished in our paradigm what specific process is impaired that is responsible for the deficits we observed. It could be a loss of memory of the correct location of the target during the delay. At least one study suggests that loss of posterior parietal lobe function by cooling impairs a monkey's ability to retain spatial information (Quintana and Fuster, 1993). Such information could be a memory of where sensory stimuli have appeared in the recent past. In this case it would then be less likely that inactivation of area LIP was responsible for the defect, as we have shown that most LIP neurons hold in memory the next planned saccade and not the location of previous stimuli (see Chapters 7 and 8). A disruption of a memorized motor plan—as would be obtained by inactivation of area LIP—, however, would also manifest itself as a loss of spatial information in the memory saccade task. The contribution of each of these processes to the deficits we observed needs to be addressed experimentally. A monkey could be trained to report the location of a previous stimulus via a response different from an eye movement or a directional movement in general, and its performance could be compared to its ability to program movements during a delay.

The fact that inactivating the PPC produces memory-saccade deficits without simultaneously lesioning other eye-movement centers such as the SC, FEF, or supplementary eye fields (SEF) establishes that the PPC (or at least some area within it, such as LIP) plays an essential role in this type of behavior. A specific deficit of memory saccade performance has been described following lesions of the FEF (Deng et al., 1986), demonstrating that this structure is also essential for the performance of memory saccades. The fact that FEF lesions also produce a specific impairment of memory saccade supports the possibility that the region inactivated in our experiments is area LIP: this area is a major source of input from the PPC to the FEF. More importantly, the fact that the PPC and the FEF are essential for

memory saccades but not for visually guided saccades in monkeys (Schiller et al., 1980) supports the neurological evidence, mentioned above, that cortical centers control saccades that are more complex than visual ones. Unlike visually guided saccades, memory saccades require that a spatial location be memorized, that a conscious plan to make a saccade be formulated, that execution of this plan be withheld for a period of time, and that a saccade to a location with no sensory stimulus be triggered by a nonspatial cue (the fixation point offset). The results of our PPC lesions and those of Deng et al. for the FEF demonstrate that at least some of these processes cannot be processed by subcortical pathways alone, but rather require processing by specific cortical eye movement centers.

Nystagmus

Single muscimol injections at a separate site produced a temporary sustained sawtooth nystagmus, with a slow phase directed contralaterally and upwards and of roughly constant velocity (though the velocity's magnitude depended on lighting conditions). Such nystagmus could be due to an imbalance of vestibular signals. Such signals have been recorded directly in two distinct regions of the parietal lobe: area 2v (Buttner and Buettner, 1978) and the parieto-insular cortex (Grüsser, 1990a, b). Vestibular signals have also been inferred in area 7a and LIP (Snyder et al., 1993). Inactivation of these regions could in principle be responsible for the effect we observed. Without anatomical confirmation, however, and because I did not test in the previous recording experiments for activity whose absence could explain the nystagmus, this hypothesis remains speculative.

Effects of multiple injections

Muscimol injections at multiple sites within the PPC resulted in a florid hemineglect syndrome. Although I only studied these symptoms using a "clinical" approach and did not quantitate them, they were obvious enough to merit reporting. Whereas in humans hemineglect is a common symptom of posterior parietal lobe lesions (for a review see

Bisiach et al., 1988), a systematic study of the effects of PPC removal in monkeys produced only extinction symptoms and no full-blown contralateral neglect syndrome (Lynch and McLaren, 1989). The effects on the monkey's reaching ability in our study confirm previous results of lesions of area 5 (disuse and clumsiness of the contralateral arm; Stein, 1976; Faugier-Grimaud et al., 1985) and area 7 (clumsiness and misreaching in the contralateral hemifield; Stein, 1976; O. Bock, R. Eckmiller, and R.A. Andersen, unpublished observations), suggesting that our multiple injections inactivated both structures.

Conclusion

The results obtained so far in this study confirm the involvement of the posterior parietal lobe of the primate in multiple processes underlying sensorimotor integration. These include the perception of spatial relationships among sensory stimuli, the allocation of attention to regions of the external world, the formation of motor plans based on sensory information, and the ability to maintain spatial representations and motor plans in memory for future use. The disruption of at least some of these processes, and perhaps of all of them, is required to obtain deficits such as hemineglect and impaired reaching and eye movements as we have seen in this study. More importantly, the specific deficit in memory saccades suggests a general direction in which to pursue future study of various cortical and subcortical structures. Many of the behavioral tasks used so far to probe nervous system function have been necessarily simplified ones, and have revealed multiple areas that are active in similar ways in relation to a given task. Introducing multiple behavioral parameters and higher level of processing requirements may open the way to clearer distinctions among the functions of various cortical areas—and perhaps help us find out some day how the brain does whatever it is that it does.

FIGURE LEGENDS

Figure 1:

Amplitudes and directions of visually guided saccades in eight directions. The target appeared at a 15 deg. eccentricity in the directions 0, 45, 90, 135, 180, 225, 270, and 315 deg. The starting position is the eye position 100 ms before the start of the saccade; all starting positions are set at the origin in these graphs. The endpoint of each vector is the eye position 100 ms after the start of the saccade. This corresponded to the final eye position after possible secondary saccades or glissades. I chose this definition of saccade endpoint instead of one relative to the end of the saccade because many saccades were slower and had irregular velocity profiles after muscimol injection, making it more difficult to identify a clear time of saccade end. In this figure and in similar ones below, the saccades made to each target are marked with a different symbol, in order to facilitate their identification in case of overlap of the endpoints of saccades to different targets. A) Visual saccades before the injection. B) Visual saccades 1 hour after muscimol injection into the PPC.

Figure 2:

Amplitudes and directions of visually guided saccades after recovery from the muscimol injection made the previous day. Conventions are as in Fig. 1.

Figure 3:

Polar plot of the amplitude of the horizontal (A) and vertical (B) components of visually guided saccades made to the 8 different targets. In this plot, as well as in all other polar plots in this chapter, the following conventions apply: 1) each direction in the polar plot corresponds to a target direction; 2) the radial position of each point in the plot is proportional to the value of the quantity being plotted; 3) squares joined by a solid line denote the values before muscimol

injection, while circles joined by a dashed line denote values 1 hour after muscimol injection into the PPC.

Figure 4:

Polar plot of latency of saccade onset (A) and peak saccadic velocity (B) for visually guided saccades before and after muscimol injection. Plot conventions are as in Fig. 3.

Figure 5:

Main sequence plot for visual saccades before and after muscimol injection. Shown is the peak saccadic velocity for each saccade in a single run of trials vs. its amplitude. Each point corresponds to one saccade. Solid squares represent values obtained before muscimol injection, while open triangles denote values obtained after injection.

Figure 6:

Amplitudes and directions of saccades to remembered locations (memory saccades) of a target appearing in one of eight different directions, before (A), and 1 hour after (B) muscimol injection into the PPC. Plot conventions are as in Fig. 1.

Figure 7:

Amplitudes and directions of memory saccades after recovery from the muscimol injection made the previous day. Conventions are as in Fig. 1.

Figure 8:

Polar plot of the amplitude of the horizontal (A) and vertical (B) components of memory saccades made to the 8 different targets, before and after muscimol injection. Conventions are as in Fig. 3.

Figure 9:

Polar plot of latency of saccade onset, relative to the fixation point's disappearance (A) and peak saccadic velocity (B) for memory saccades before and after muscimol injection. Plot conventions are as in Fig. 3.

Figure 10:

Main sequence plot for memory saccades before and after muscimol injection. Shown is the peak saccadic velocity for each saccade in a single run of trials vs. its amplitude. Plot conventions are as in Fig. 5.

Figure 11:

Example of the nystagmus observed after muscimol injection at a site c. 2.0 mm away from the site where muscimol injection produced the saccade deficits shown in Figs. 1-10. Shown here are traces of horizontal (e_x) and vertical (e_y) eye position vs. time (100 ms/division) for a segment of one trial while the monkey attempted to maintain fixation.

Figure 12:

Plot of horizontal vs vertical eye position throughout a series of trials in which the monkey had to maintain fixation at each of the locations marked by crosses (not all crosses are visible; they are regularly spaced, however, at 20 deg. horizontal and vertical intervals). Fig. 11 shows a segment of one of these trials over time. Eye position was sampled at 500 Hz.

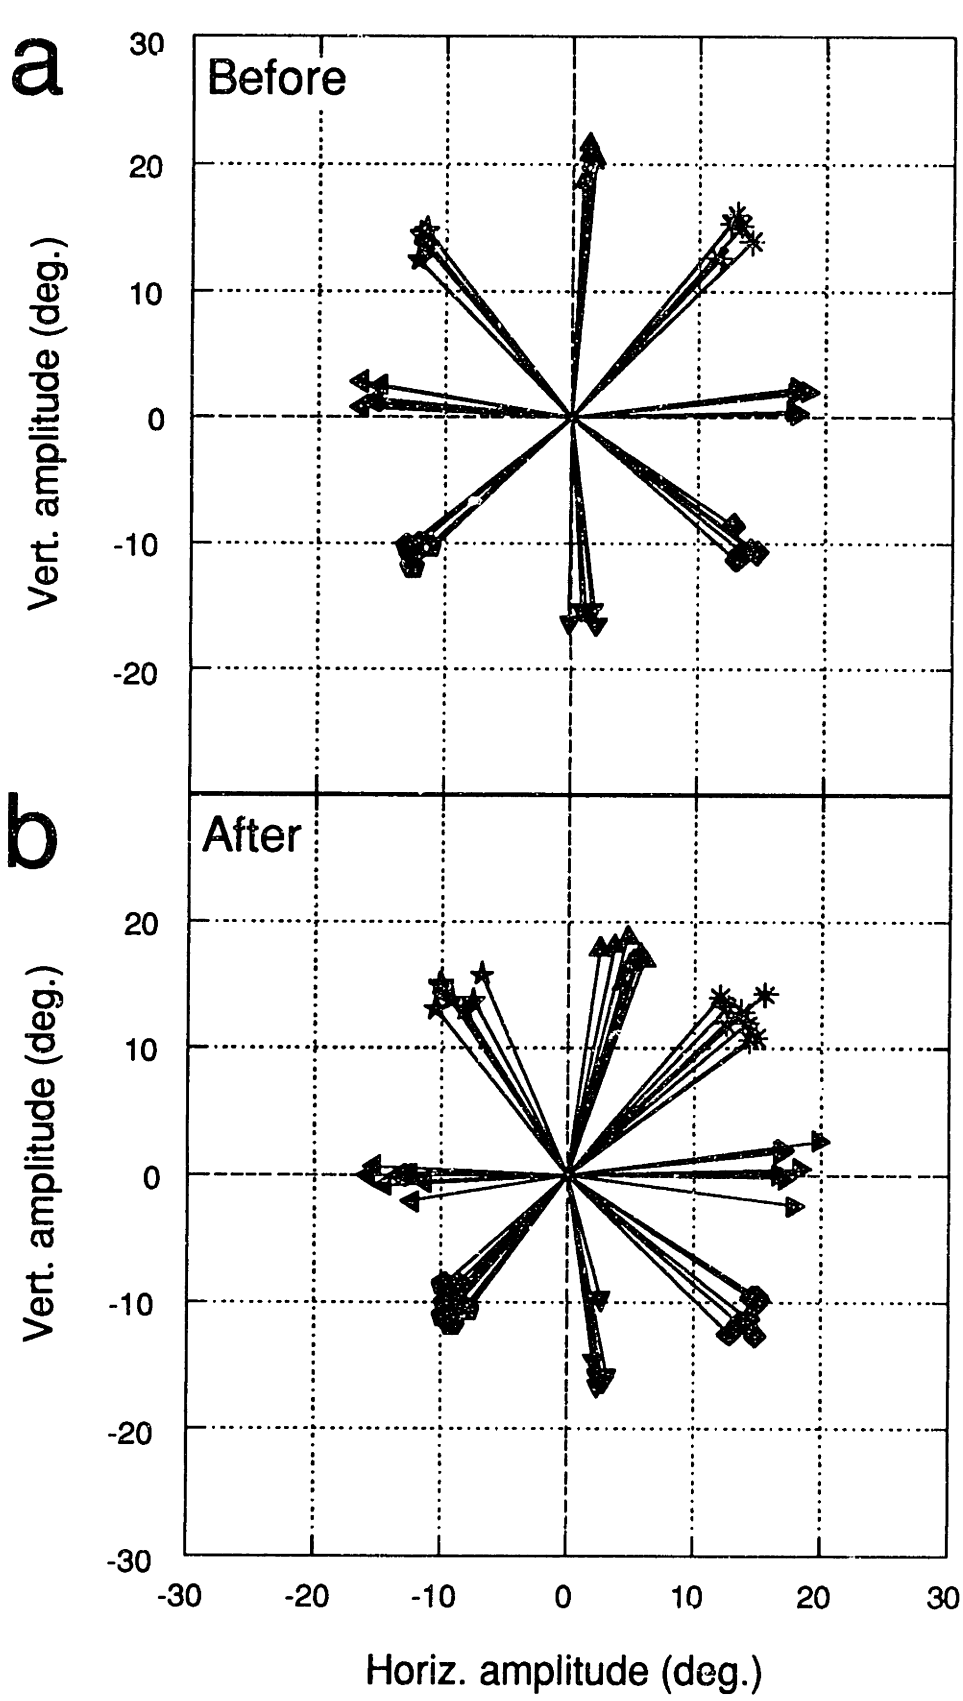
Visual Saccades
a
Before
b
After
Vert. amplitude (deg.)
Vert. amplitude (deg.)
Horiz. amplitude (deg.)
30
20
10
0
-10
-20
-30
-30
-20
-10
0
10
20
30

Figure 1

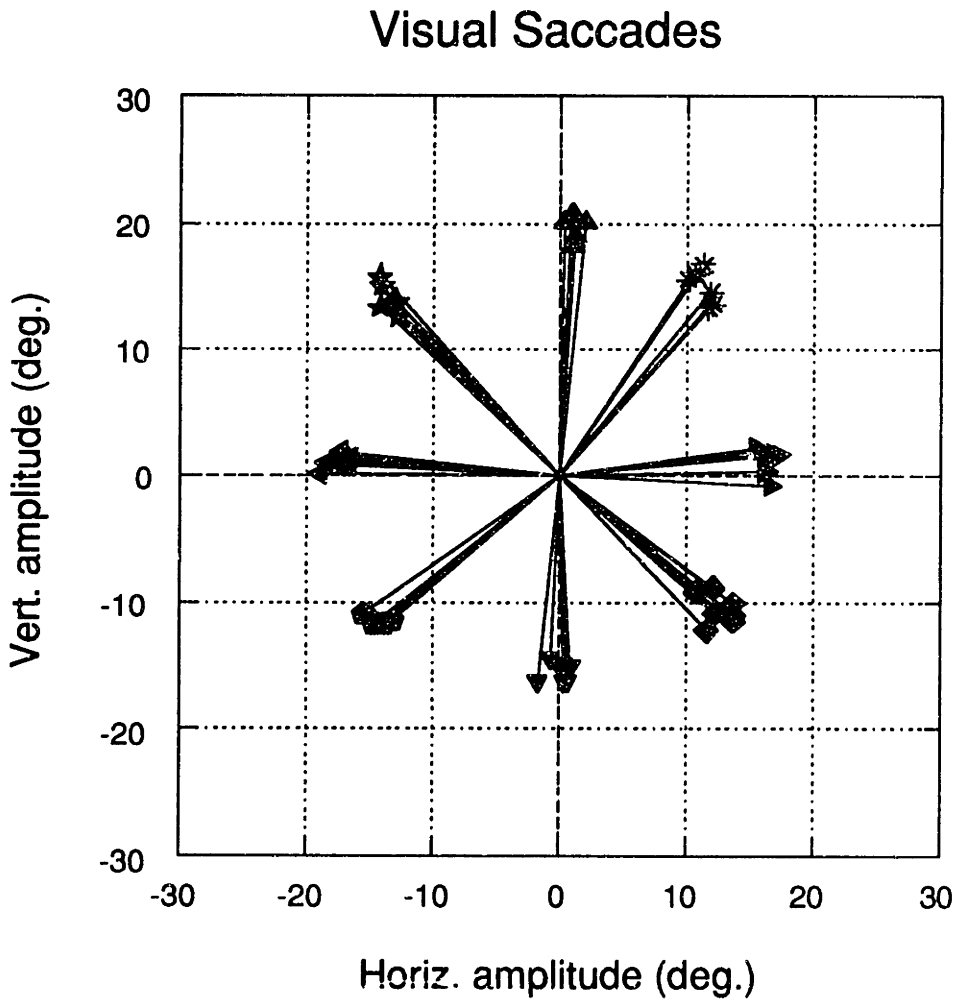
Visual Saccades
30
20
10
0
-10
-20
-30
Vert. amplitude (deg.)
-30
-20
-10
0
10
20
30
Horiz. amplitude (deg.)

Figure 2

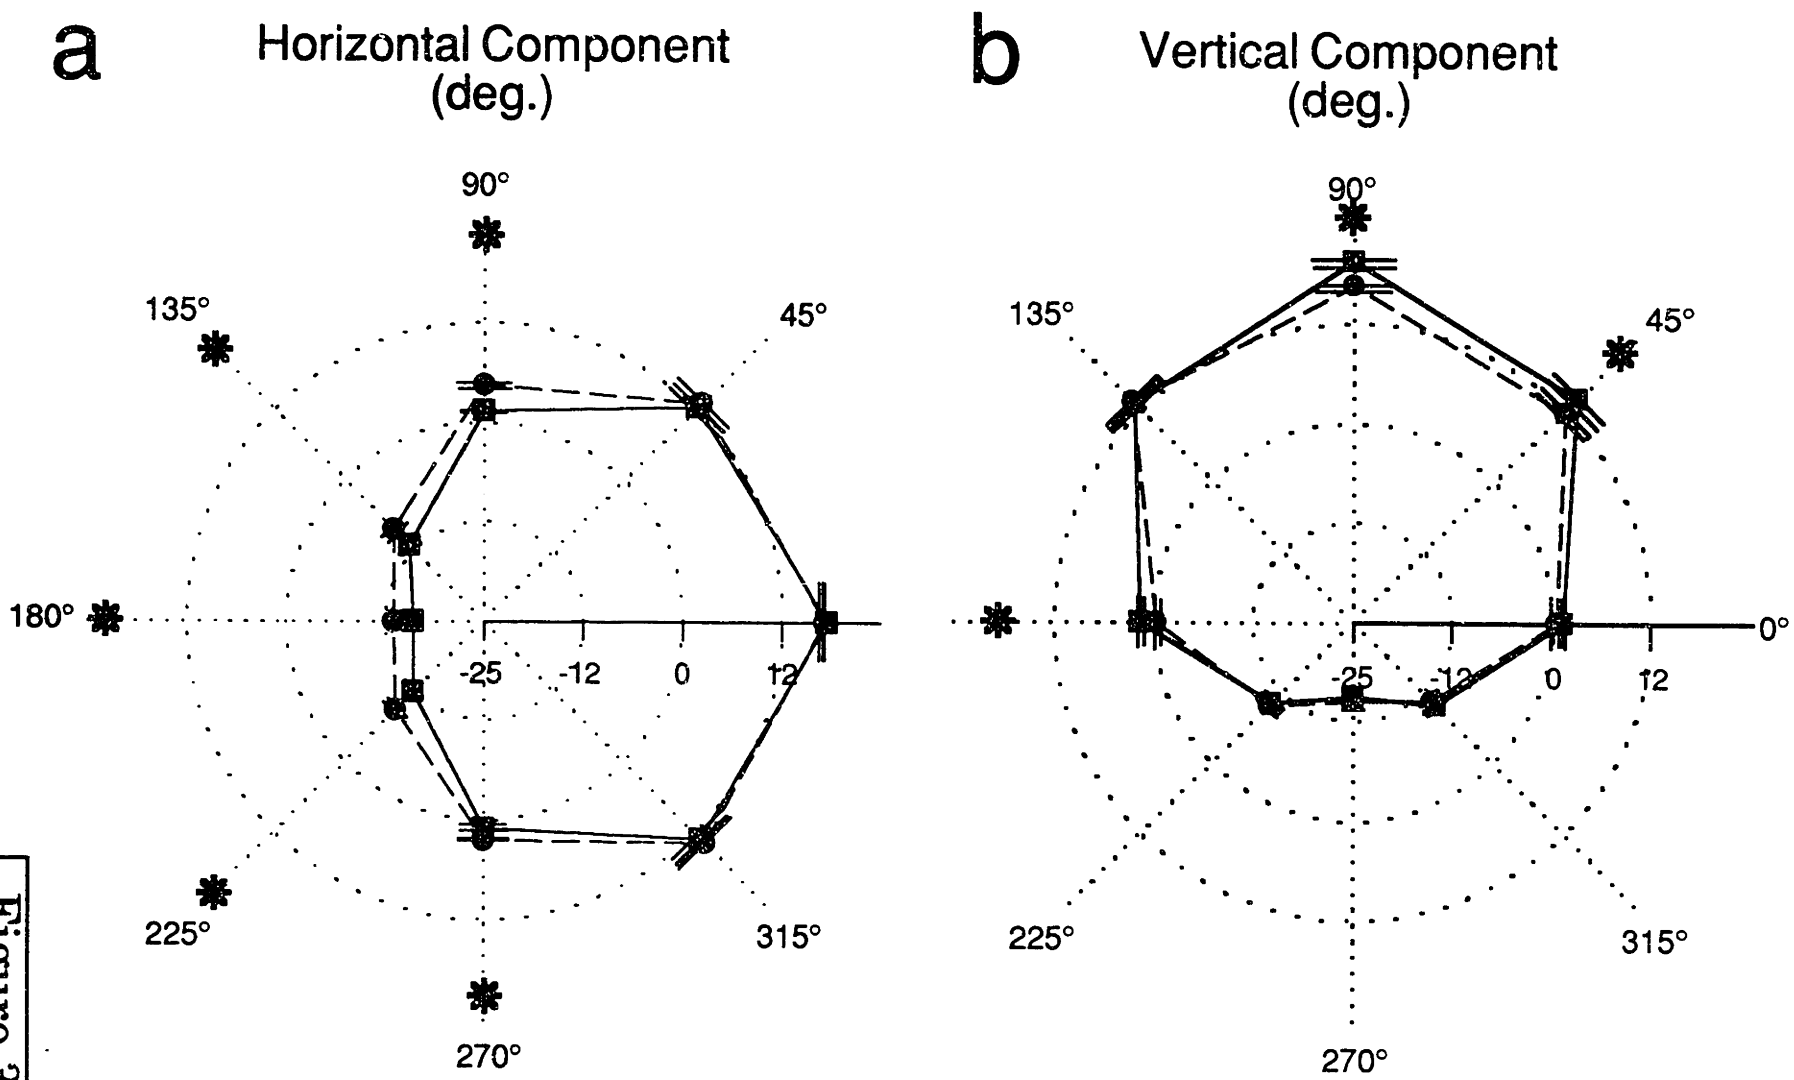
a
Horizontal Component
(deg.)
90°
135°
45°
180°
225°
315°
270°
-25
-12
0
12
b
Vertical Component
(deg.)
90°
135°
45°
0°
225°
315°
270°
-25
-12
0
12

Figure 3

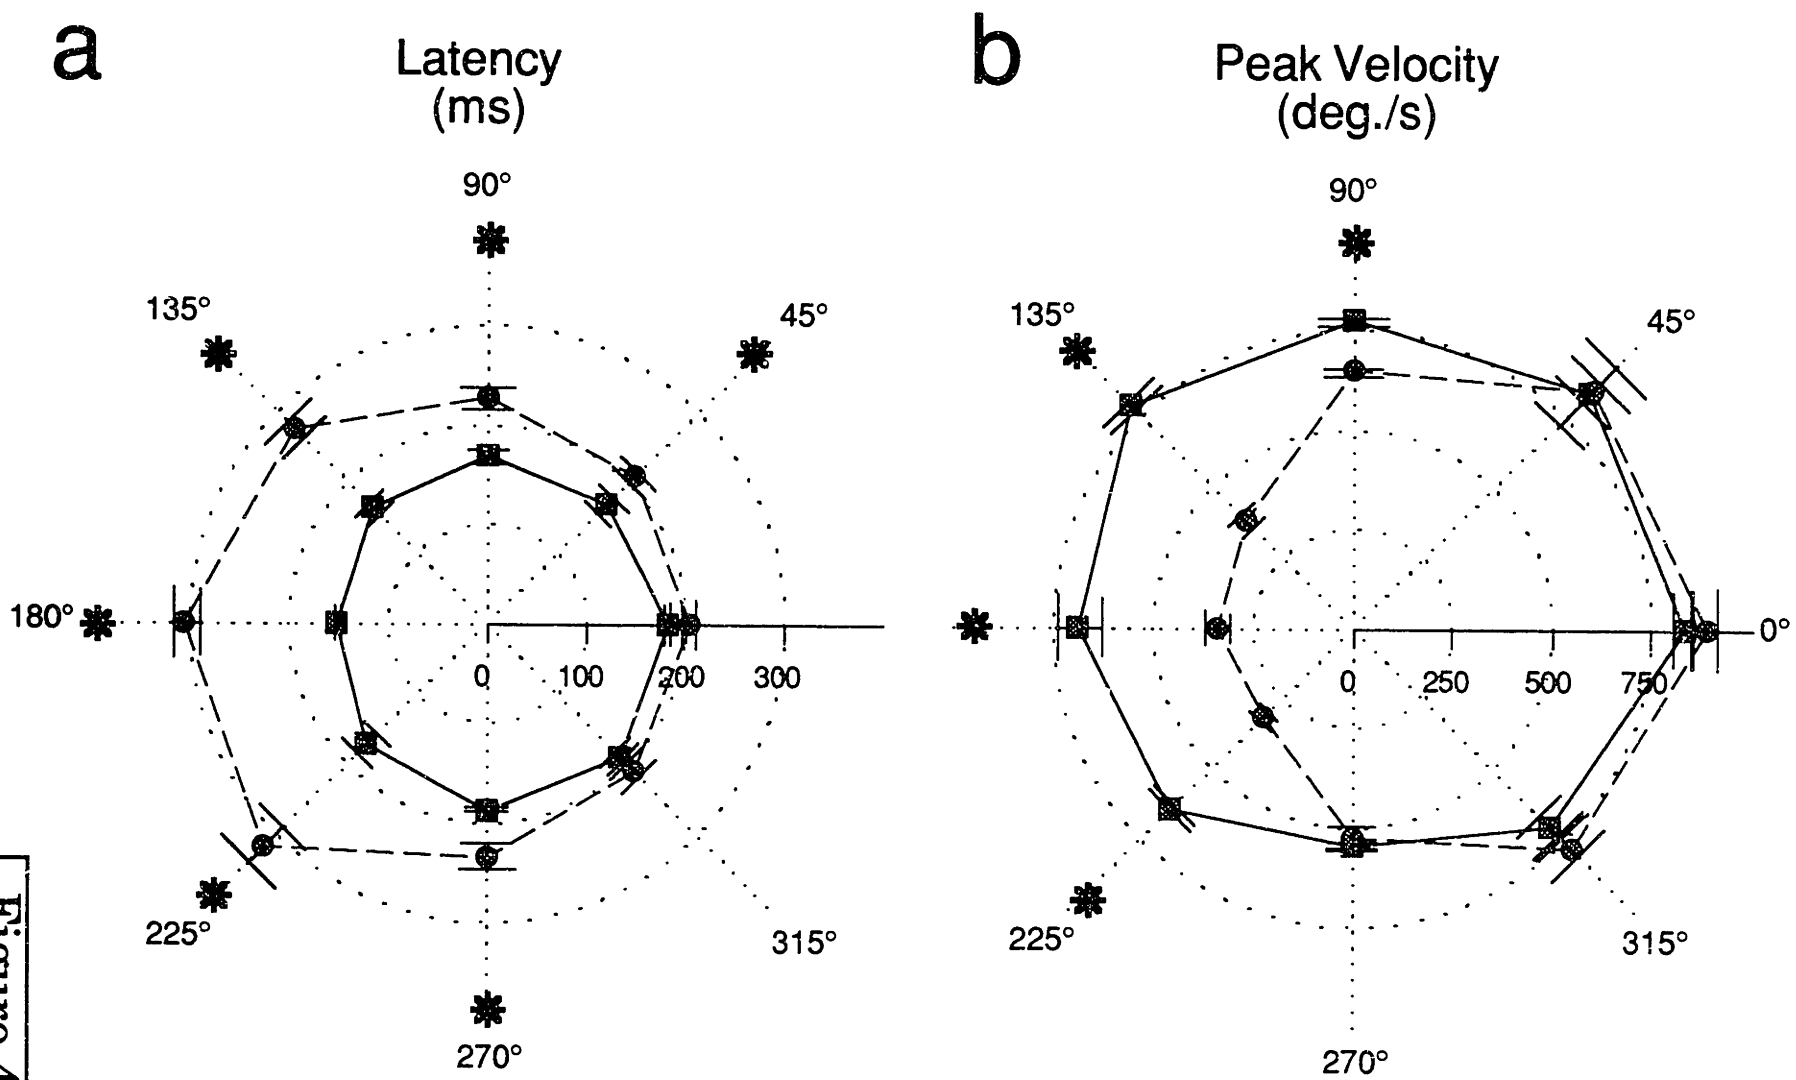
a
Latency
(ms)
90°
135°
45°
180°
0°
225°
315°
270°
0
100
200
300
b
Peak Velocity
(deg./s)
90°
135°
45°
0°
225°
315°
270°
0
250
500
750

Figure 4

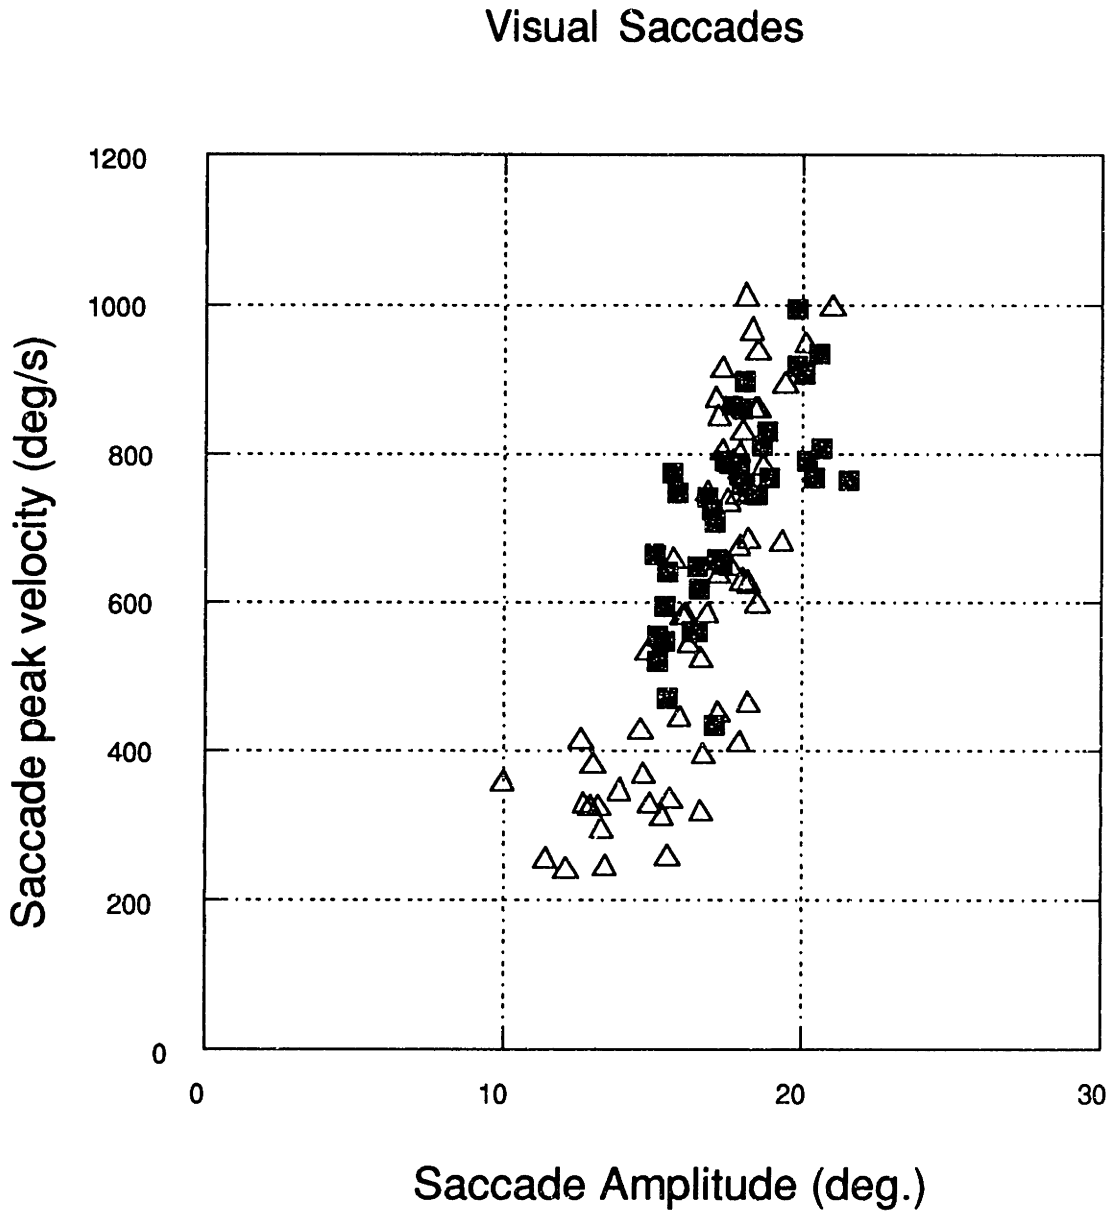
Visual Saccades
Saccade peak velocity (deg/s)
1200
1000
800
600
400
200
0
0
10
20
30
Saccade Amplitude (deg.)

Figure 5

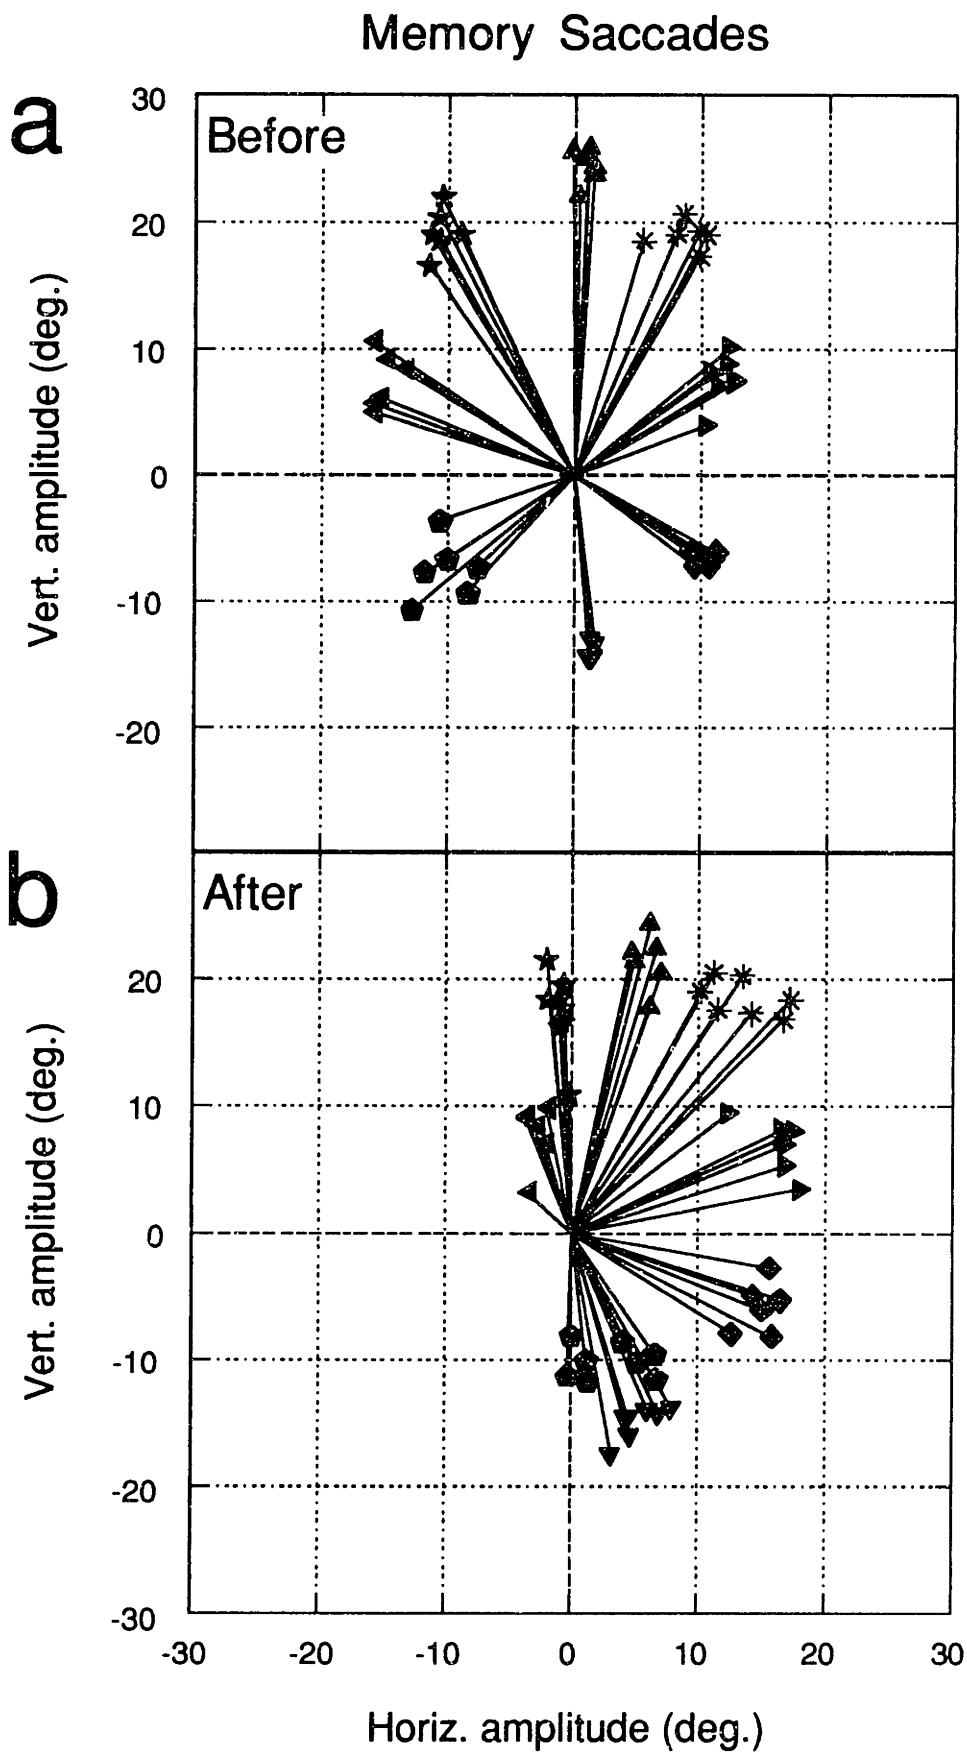
Memory Saccades
a
Before
b
After
Vert. amplitude (deg.)
Vert. amplitude (deg.)
Horiz. amplitude (deg.)
30
20
10
0
-10
-20
-30
-30
-20
-10
0
10
20
30

Figure 6

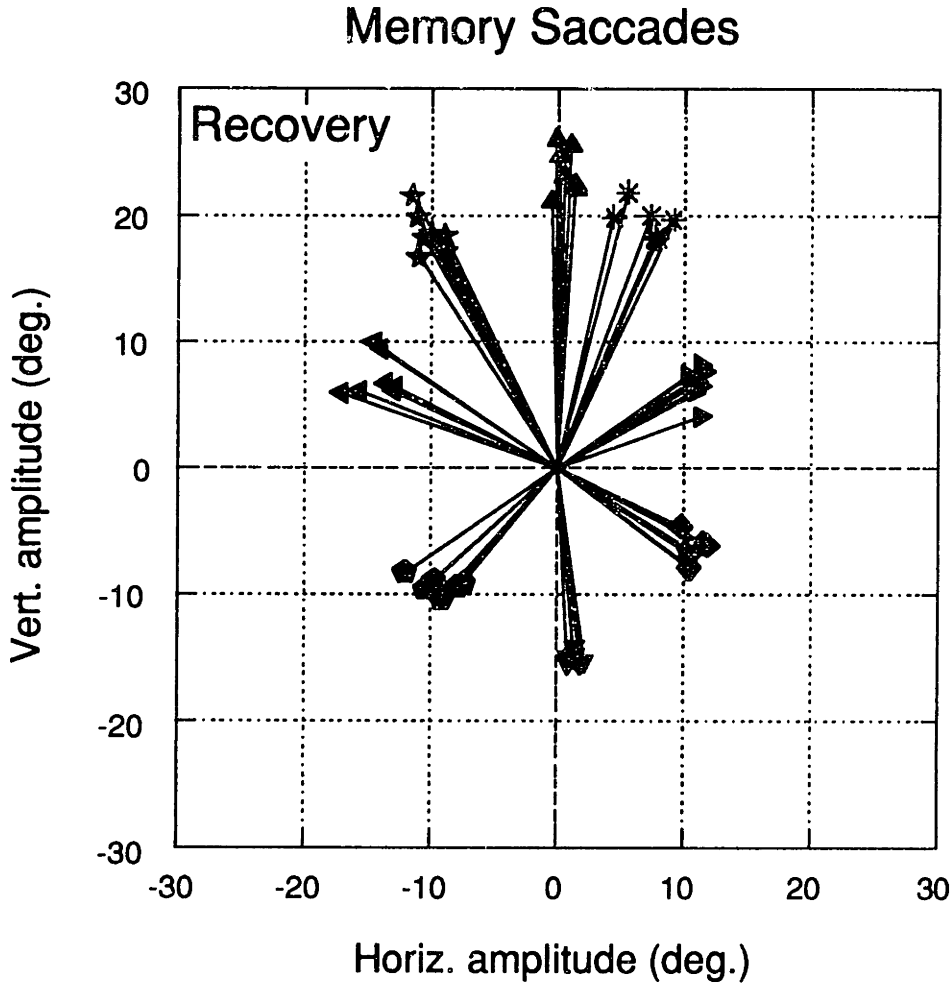
Memory Saccades
Recovery
30
20
10
0
-10
-20
-30
-30
-20
-10
0
10
20
30
Vert. amplitude (deg.)
Horiz. amplitude (deg.)

Figure 7

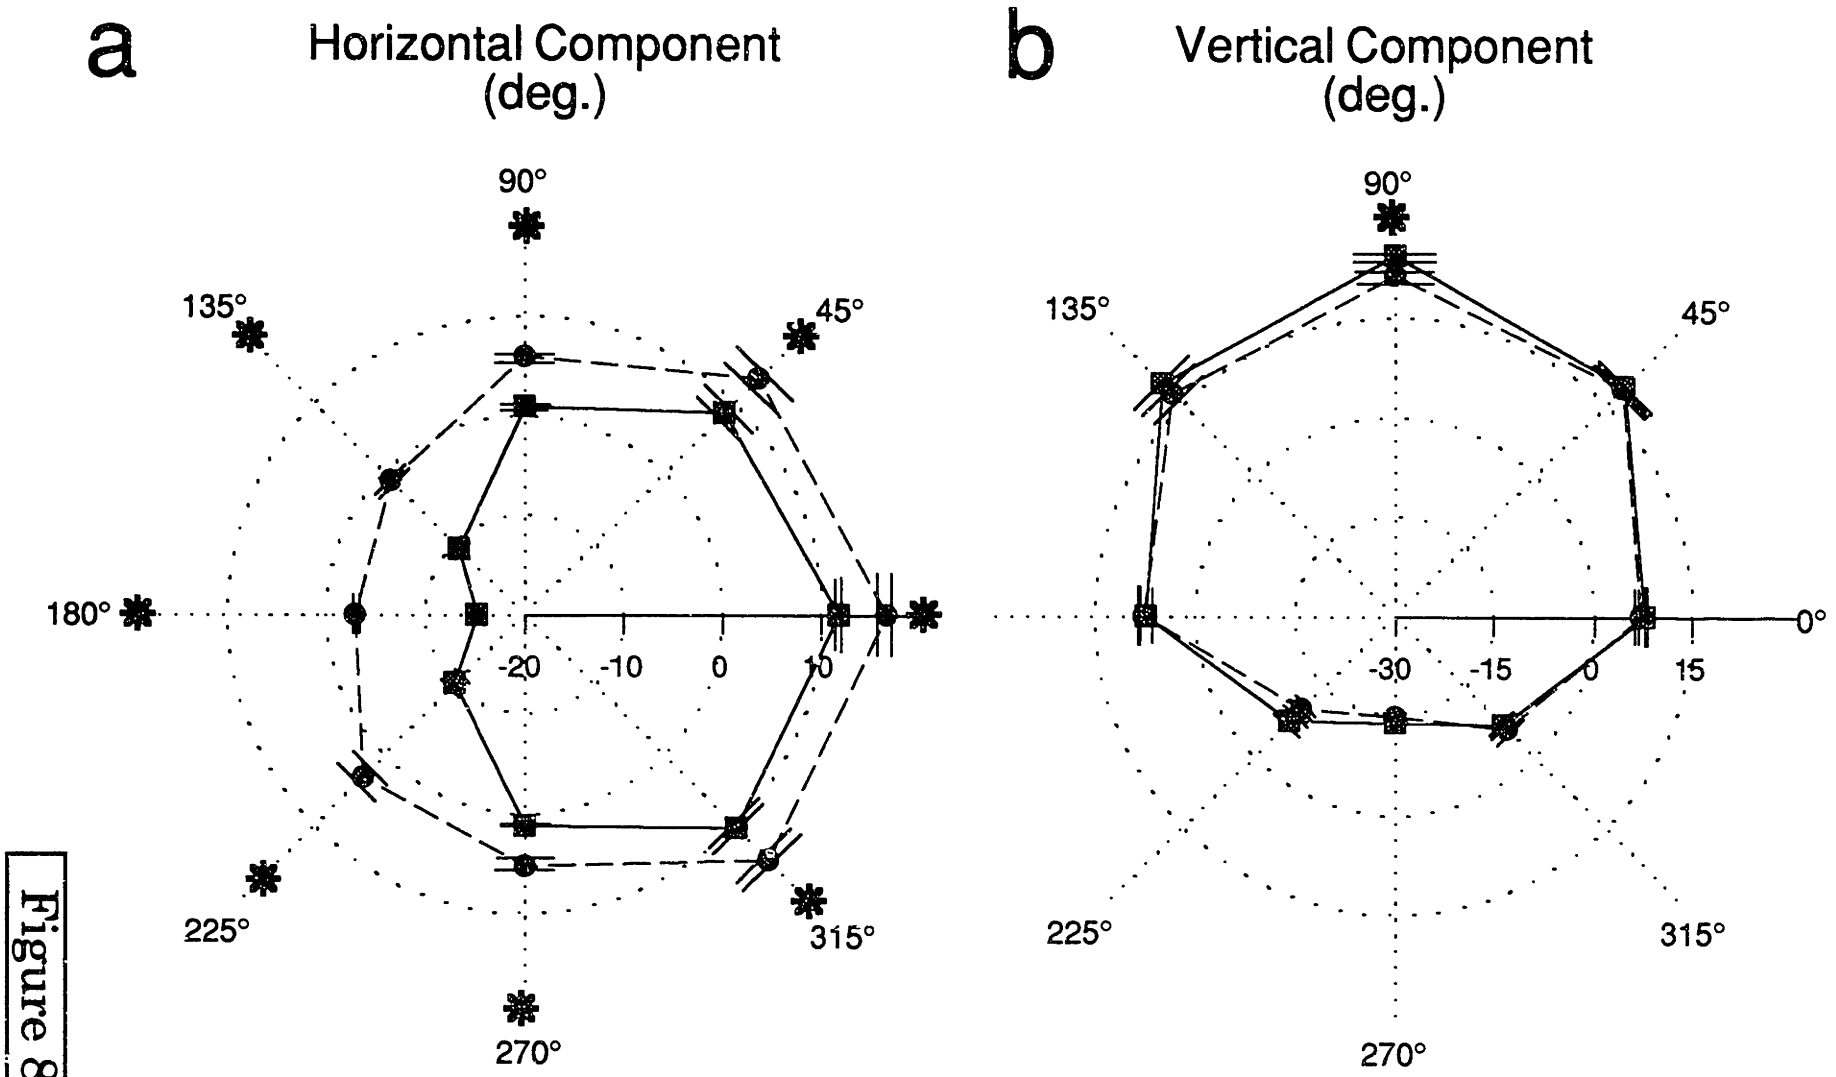
a
Horizontal Component
(deg.)
90°
135°
45°
180°
225°
315°
270°
-20
-10
0
10
b
Vertical Component
(deg.)
90°
135°
45°
0°
225°
315°
270°
-30
-15
0
15

Figure 8

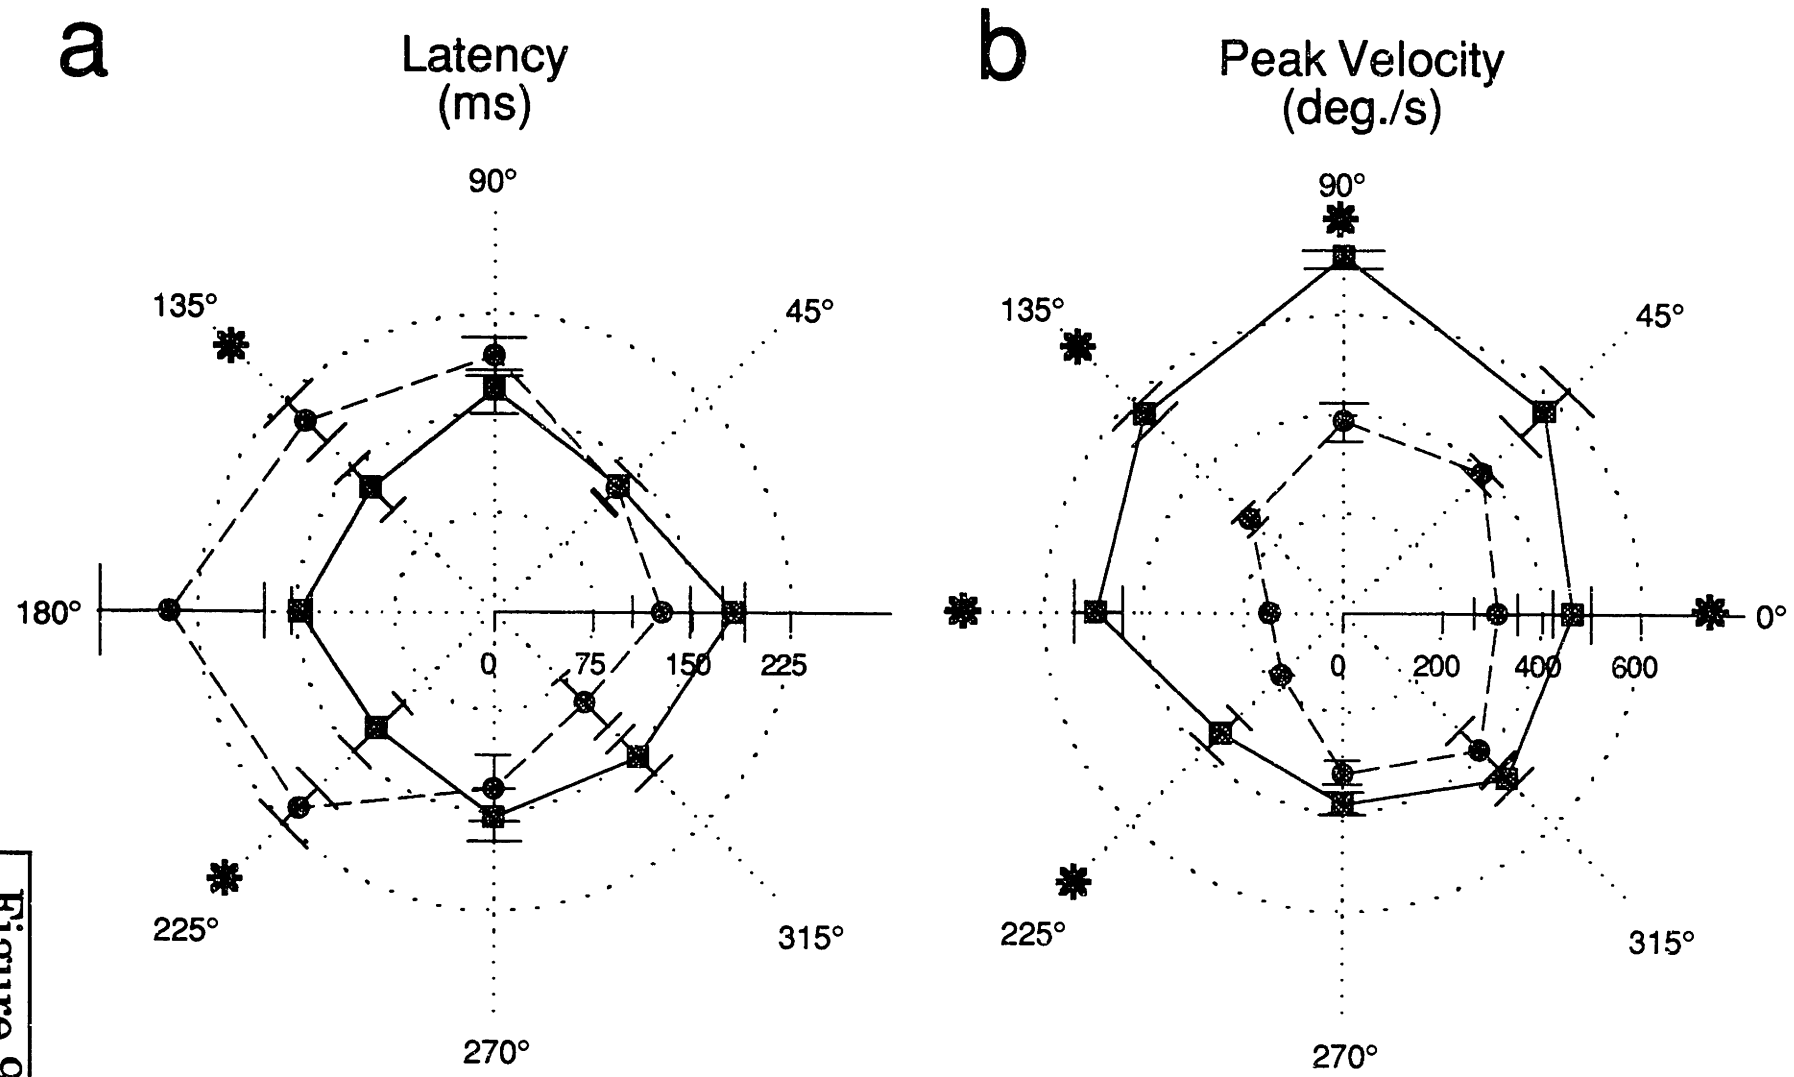
a
Latency
(ms)
90°
135°
45°
180°
0°
225°
315°
270°
0
75
150
225
b
Peak Velocity
(deg./s)
90°
135°
45°
0°
225°
315°
270°
0
200
400
600

Figure 9

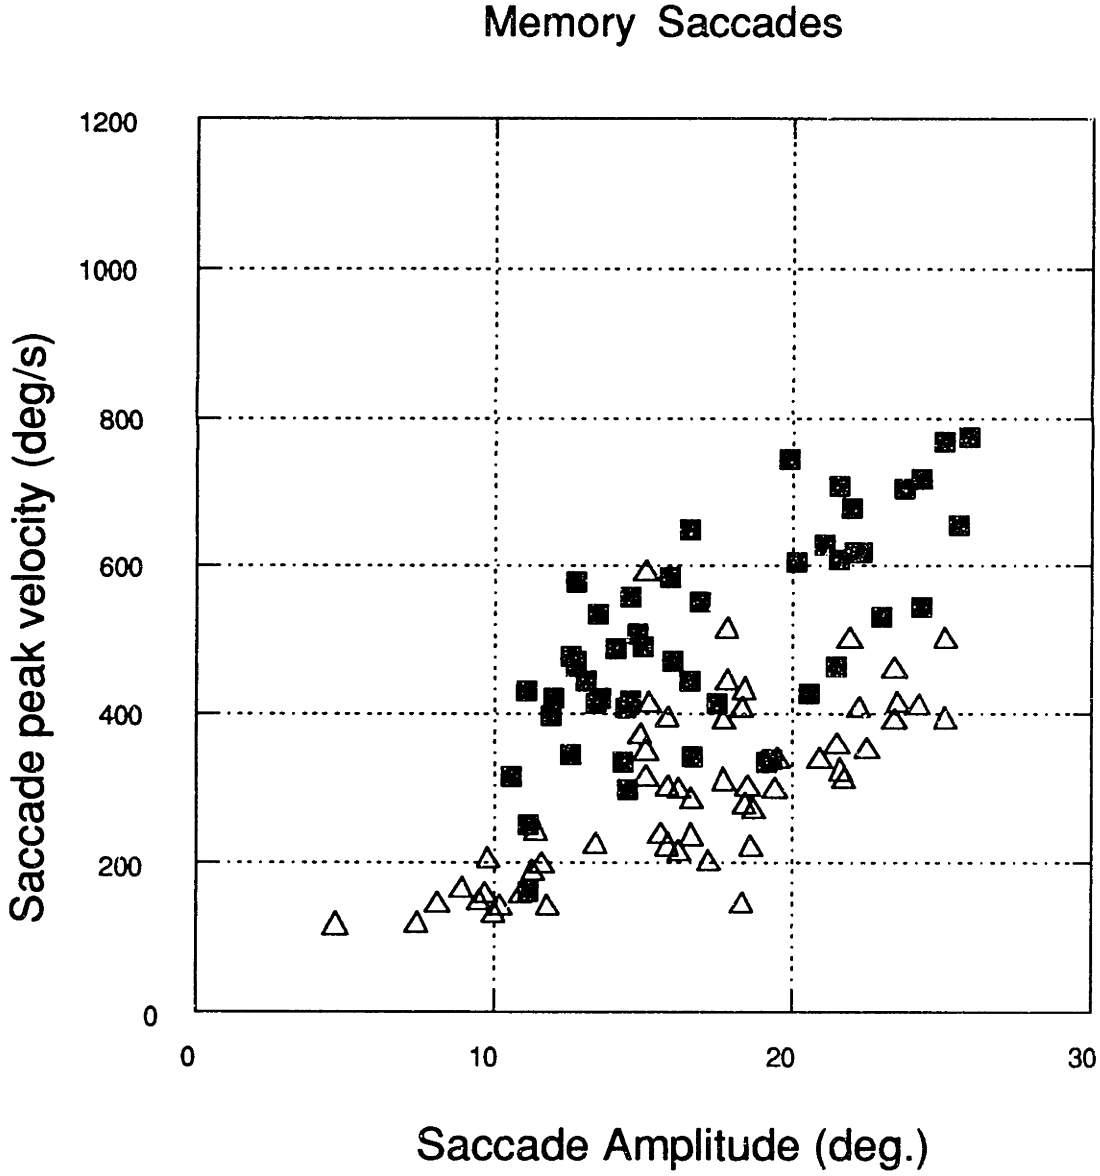
Memory Saccades
Saccade peak velocity (deg/s)
1200
1000
800
600
400
200
0
0
10
20
30
Saccade Amplitude (deg.)

Figure 10

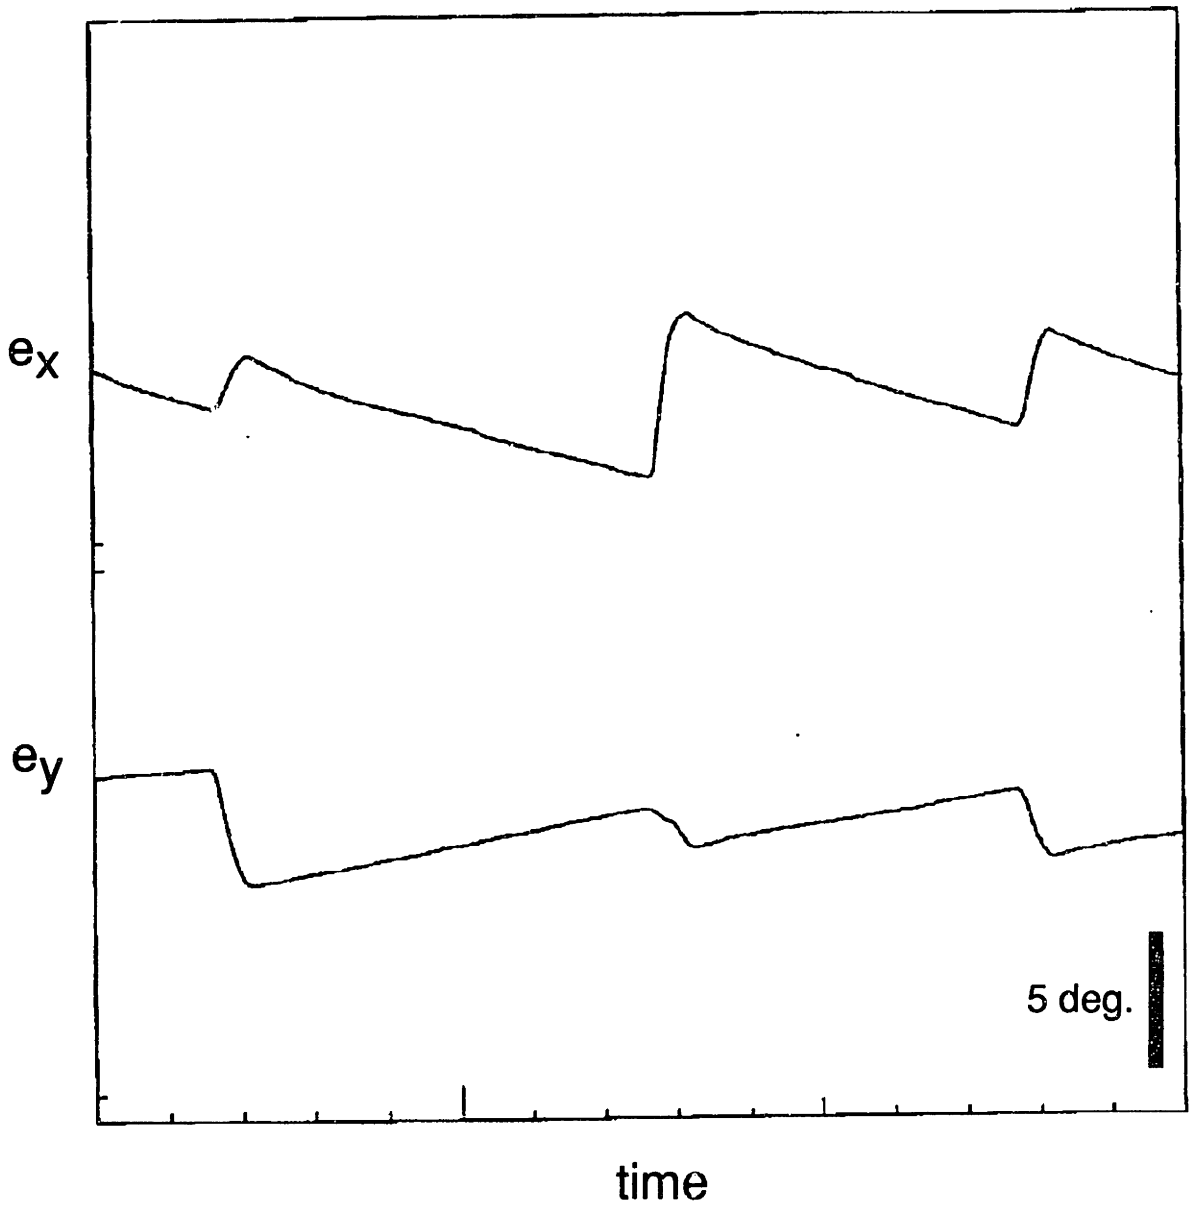

Figure 11

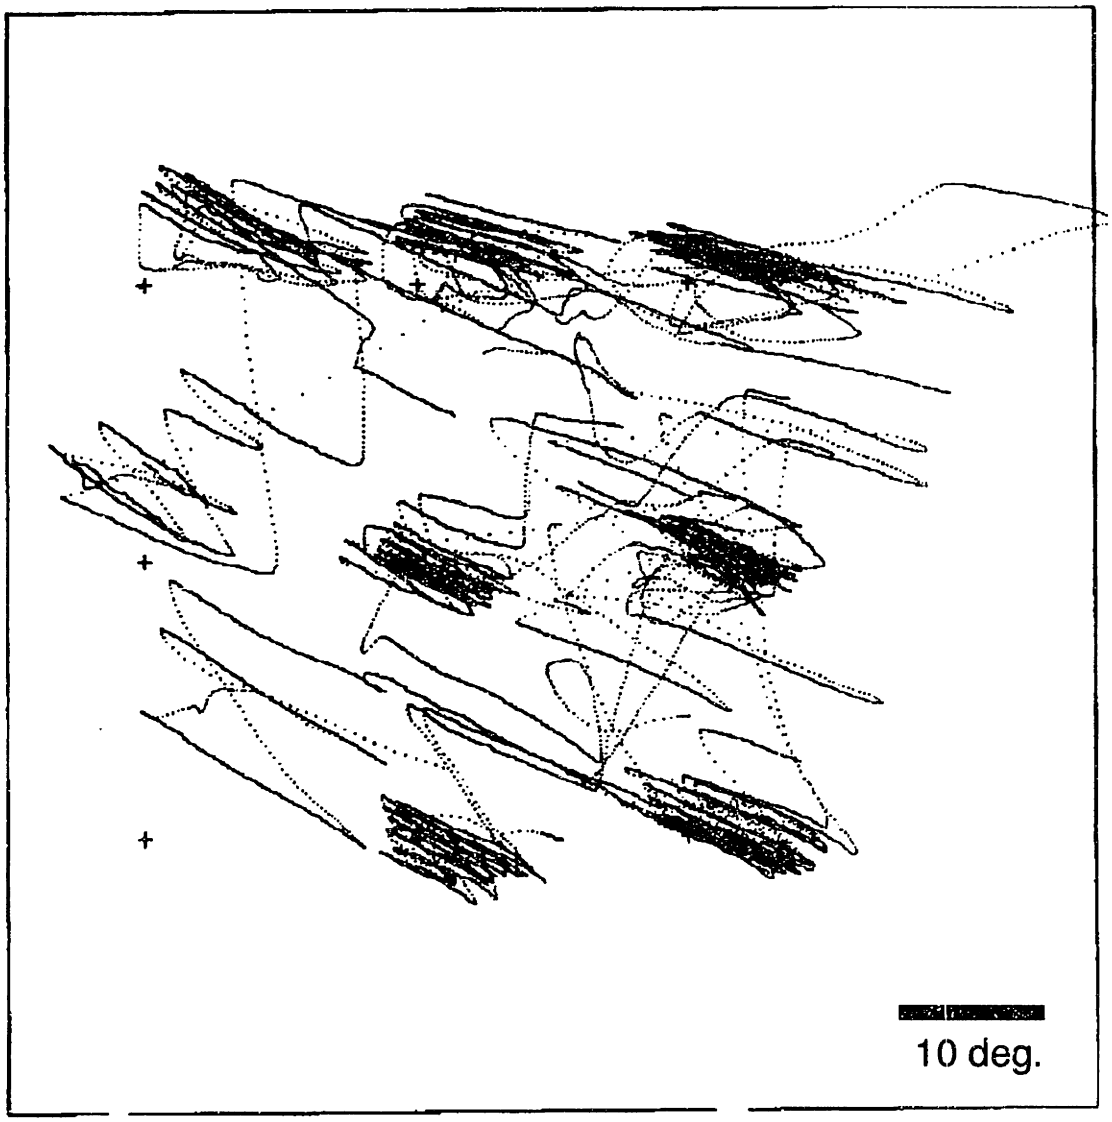

e_x

e_y

Figure 12

Bibliography

Albano, J. E., Mishkin, M., Westbrook, L. E., and Wurtz, R. H. (1982) Visuomotor deficits following ablation of monkey superior colliculus. *J Neurophysiol* 48(2):338-51.

Alsternack, B., Lundberg, A., Norsell, U., and Sybirski, E. (1981) Integration on descending motor pathways controlling the forelimb in the cat. IX. Differential behavioral effects after spinal cord lesions interrupting defined pathways from higher centres to motoneurons. *Exp Brain Res* 42:299-318.

Altman, J. A., Balonov, L. J., and Deglin, V. L. (1979) Effects of unilateral disorder of the brain hemisphere. *Neuropsychol* 17:295-301.

Andersen, R. A., and Mountcastle, V. B. (1983) The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex. *J Neurosci* 3(3):532-548.

Andersen, R. A., Asanuma, C., and Cowan, W. M. (1985a) Callosal and prefrontal associational projecting cell populations in area 7a of the macaque monkey: a study using retrogradely transported fluorescent dyes. *J Comp Neurol* 232:443-455.

Andersen, R. A., Essick, G. K., and Siegel, R. M. (1985b) Encoding of spatial locations by posterior parietal neurons. *Science* 230:456-458.

Andersen, R. A., Essick, G. K., and Siegel, R. M. (1987) Neurons of area 7a activated by both visual stimuli and oculomotor behavior. *Exp Brain Res* 67:316-322.

Andersen, R. A. (1987) The role of the inferior parietal lobule in spatial perception and visual-motor integration. In *The handbook of physiology. Section 1: The nervous system, Vol. IV. Higher functions of the brain, Part 2*, F. Plum, V. B. Mountcastle, and S. T. Geiger, eds., vol. IV, pp. 483-518, American Physiological Society, Bethesda, MD.

Anderson, C. H., and Van Essen, D. C. (1987) Shifter circuits: a computational strategy for dynamic aspects of visual processing. *Proc Natl Acad Sci USA* 84:6297-6301.

Andersen, R. A., and Zipser, D. (1988) The role of the posterior parietal cortex in coordinate transformations for visual-motor integration. *Can J Physiol Pharmacol* 66(4):488-501.

Anderson, J. A., and Rosenfeld, E. (eds.) (1988) *Neurocomputing*, MIT Press, Cambridge.

Andersen, R. A. (1989) Visual and eye movement functions of the posterior parietal cortex. *Annu Rev Neurosci* 12:377-403.

Andersen, R. A., and Gnadt, J. W. (1989) Posterior parietal cortex. In *The neurobiology of saccadic eye movements*, R. H. Wurtz, and M. E. Goldberg, eds., pp. 315-335, Elsevier Science Publishers BV, Amsterdam.

Andersen, R. A., Asanuma, C., Essick, G., and Siegel, R. M. (1990a) Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J Comp Neurol* 296(1):65-113.

Andersen, R. A., Bracewell, R. M., Barash, S., Gnadt, J. W., and Fogassi, L. (1990b) Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J Neurosci* 10(4):1176-96.

Andersen, R. A., Brotchie, P. R., and Mazzoni, P. (1992) Evidence for the lateral intraparietal area as the parietal eye field. *Curr Opin Neurobiol* 2:840-846.

Andrews, P. R., and Johnson, G. A. R. (1979) GABA agonists and antagonists. *Biochem Pharmacol* 28:2697-2702.

Auerbach, C., and Sperling, P. (1974) A common auditory-visual space: evidence for its reality. *Percept Psychophys* 16(1):129-135.

Baizer, J. S., Ungerleider, L. G., and Desimone, R. (1991) Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *J Neurosci* 11(1):168-190.

Balint, R. (1909) Seelenlähmung des "Schauens", optische Ataxie, räumliche Störung der Aumferksamkeit. *Monatsschr. Psychiatr. Neurol.* 25:51-81.

Barash, S., Bracewell, R. M., Fogassi, L., and Andersen, R. A. (1989) Interactions of visual and motor-planning activitites in the lateral intra-parietal area (LIP). *Soc Neurosci Abst* 2:1203.

Barash, S., Bracewell, R. M., Fogassi, L., Gnadt, J. W., and Andersen, R. A. (1991a) Saccade-related activity in the lateral intraparietal area. I. Temporal properties; comparison with area 7a. *J Neurophysiol* 66(3):1095-1108.

Barash, S., Bracewell, R. M., Fogassi, L., Gnadt, J. W., and Andersen, R. A. (1991b) Saccade-related activity in the lateral intraparietal area. II. Spatial properties. *J Neurophysiol* 66(3):1109-1124.

Barto, A. G. (1985) Learning by statistical cooperation of self-interested neuron-like computing elements. *Hum Neurobiol* 4:229-256.

Barto, A. G., and Anandan, P. (1985) Pattern-recognizing learning stochastic automata. *IEEE Trans Sys, Man, Cybern* 15(3):360-375.

Barto, A. G., and Jordan, M. I. (1987) Gradient following without backpropagation in layered networks. *Proceedings of IEEE International Conference on Neural Networks*, San Diego, CA, June 1987, vol. II, 629-636.

Barto, A. G. (1989) From chemotaxis to cooperativity: abstract exercises in neuronal learning strategies. In *The computing neuron*, R. M. Durbin, R. C. Miall, and G. J. Mitchison, eds., pp. 73-98, Addison-Wesley, New York.

Bender, M. B., and Diamond, S. P. (1965) An analysis of auditory perceptual defects with observations on the localization of dysfunction. *Brain* 88:675-686.

Bisiach, E., and Luzzatti, C. (1978) Unilateral neglect of representational space. *Cortex* 14:129-133.

Bisiach, E., Luzzatti, C., and Perani, D. (1979) Unilateral neglect, representational schema and consciousness. *Brain* 102:757-765.

Bisiach, E., Cornacchia, L., Sterzi, R., and Vallar, G. (1984) Disorders of perceived auditory lateralization after lesions of the right hemisphere. *Brain* 107(1):37-52.

Bisiach, E., and Vallar, G. (1988) Hemineglect in humans. In *Handbook of neuropsychology*, F. Boller, and J. Grafman, eds., vol. 1, pp. 195-222, Elsevier, New York.

Blatt, G. J., Andersen, R. A., and Stoner, G. R. (1990) Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *J Comp Neurol* 299(4):421-45.

Blauert, J. (1983) *Spatial hearing: the psychophysics of human sound localization*. MIT Press, Cambridge, MA.

Bracewell, R. M. (1991). *On the posterior parietal cortex and saccadic eye movements*. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA.

Bracewell, R. M., Barash, S., Mazzoni, P., and Andersen, R. A. (1991) Neurones in the macaque lateral intraparietal cortex (LIP) appear to encode the next intended saccade. *Soc Neurosci Abst* 17:1282.

Brain, W. R. (1941) Visual disorientation with special reference to lesions of the right cerebral hemisphere. *Brain* 64:244-272.

Brotchie, P. R., and Andersen, R. A. (1991) A body centred coordinate system in posterior parietal cortex. *Soc Neurosci Abst* 17:1281.

Brown, T. H., Kairiss, E. W., and Keenan, C. L. (1990) Hebbian synapses: biophysical mechanisms and algorithms. *Annu Rev Neurosci* 13:475-511.

Bruce, C. J., Desimone, R., and Gross, C. G. (1981) Visual properties of neurons in a polysensory area in the superior temporal sulcus of the macaque. *J Neurophysiol* 46:369-384.

Bushnell, M. C., Goldberg, M. E., and Robinson, D. L. (1981) Behavioral enhancement of visual responses in monkey cerebral cortex. I. Modulation in posterior parietal cortex related to selective visual attention. *J Neurophysiol* 46:755-772.

Butter, U., and Buettner, U. W. (1978) Parietal cortex (2v) neuronal activity in the alert monkey during natural vestibular and optokinetic stimulation. *Brain Res* 153:392-397.

Cheney, P. D., and Fetz, E. E. (1980) Comparable patterns of muscle facilitation evoked by individual corticomotoneuronal (CM) cells and by cinlge intracortical microstimuli in primates: evidence for functional groups of CM cells. *J Neurophysiol* 53:786-804.

Cho, S., and Reggia, J. A. (1992) Learning visual coordinate transformations with competition. *Proceedings of the International Joint Conference On Neural Networks*, Baltimore, MD, USA, June 1992, vol. IV, 49-54.

Cottrell, G. W., Munro, P., Zipser, D. (1987) *Image compression by back-propagation: an example of extensional programming.*ICS Report 8702, La Jolla, CA 92093.

Critchley, M. (1953) *The parietal lobes*. Arnold, London.

De Renzi, E. (1982) *Disorders of space exploration and cognition*. Wiley, Chichester.

De Renzi, E., Gentilini, M., and Pattacini, F. (1984) Auditory extinction following hemisphere damage. *Neuropsychol* 22:733-744.

Deng, S. -., Goldberg, M. E., Segraves, M. A., Ungerleider, L. G., and Mishkin, M. (1986) The effect of unilateral ablation of the frontal eye fields on saccadic performance in the monkey. In *Adaptive Processes in Visual and Oculomotor Systems*, E. L. Keller, and D. S. Zee, eds., pp. 201-208, Pergamon, Oxford.

Dominey, P. F., and Arbib, M. A. (1992) A cortical-subcortical model for generation of spatially accurate sequential saccades. *Cerebral Cortex* 2(2):153-175.

Droulez, J., and Berthoz, A. (1991) A neural network model of sensoritopic maps with predictive short-term memory properties. *Proc Natl Acad Sci USA* 88:9653-9657.

Duhamel, J. -., Colby, C. L., and Goldberg, M. E. (1992a) The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255:90-92.

Duhamel, J. -., Goldberg, M. E., Fitzgibbon, E. J., Sirigu, A., and Grafman, J. (1992b) Saccadic dysmetria in a patient with a right frontoparietal lesion: the importance of corollary discharge for accurate spatial behavior. *Brain* 115:1387-1402.

Efron, R., Crandall, P. H., Koss, B., Divenyi, P. L., and Yund, E. W. (1983) Central auditory processing. III. The "cocktail party" effect and anterior temporal lobectomy. *Brain and Language* 19:254-263.

Evarts, E. V., and Tanji, K. (1974) Gating of motor cortex reflexes by prior instruction. *Brain Res* 71:479-494.

Evarts, E. V. (1981) Role of motor cortex in voluntary movements in primates. In *Handbook of Physiology: Section 1. Volume 2. The Nervous System: Motor Control*, V. B. Brooks, ed., pp. 1083-1120, Am Phys Soc, Bethesda.

Evarts, E. V. (1986) Motor cortex output in primates. In *Cerebral Cortex Vol. 5*, E. G. Jones, and A. Peters, eds., Plenum, NY.

Farah, M. J., Wong, A. B., Monheit, M. A., and Morrow, L. A. (1989) Parietal lobe mechanisms of spatial attention: modality-specific or supramodal? *Neuropsychol* 27(4):461-470.

Faugier-Grimaud, S., Frenois, C., and Stein, D. G. (1978) Effects of posterior parietal lesions on visually guided behavior in monkeys. *Neuropsychol* 16:151-168.

Ferrier, D. (1876) *Functions of the brain.* Smith, Elder, London.

Funahashi, S., Bruce, C. J., and Goldman-Rakic, P. S. (1993) Dorsolateral prefrontal lesions and oculomotor delayed-response performance: evidence for mnemonic "scotomas". *J Neurosci* 13:1479-1497.

Georgopoulos, A., Schwartz, A., and Kettner, R. (1986) Neuronal population coding of movement direction. *Science* 233:1416-1419.

Georgopoulos, A. P., Kettner, R. E., and Schwartz, A. B. (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *J Neurosci* 8(8):2928-37.

Georgopoulos, A. P., and Grillner, S. (1989) Visuomotor coordination in reaching and locomotion. *Science* 245(4923):1209-10.

Georgopoulos, A. P. (1991) Higher order motor control. *Annu Rev Neurosci* 14:361-77.

Gnadt, J. W., and Andersen, R. A. (1988) Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70:216-220.

Godschalk, M., Lemon, R. N., Kuypers, H. G. J., and van der Steen, J. (1985) The involvement of monkey premotor cortex neurones in preparation of visually cues arm movements. *Behav Brain Res* 18:143-57.

Goldberg, M. E., and Robinson, D. L. (1977) Visual responses of neurons in monkey inferior parietal lobule: the physiologic substrate of attention and neglect. *Neurol* 27:350.

Goldberg, G. (1985) Supplementary motor area structure and function: Review and hypotheses. *Behav Brain Sci* 8:567-616.

Goldberg, M. E., and Segraves, M. A. (1989) The visual and frontal cortices. *The neurobiology of saccadic eye movements*, R. H. Wurtz, and M. E. Goldberg, eds., pp. 283-313, Elsevier Science Publishers BV, Amsterdam.

Goldberg, M. E., Colby, C. L., and Duhamel, J. R. (1989) Parietal visual neurons achieve spatial accuracy by coordinate transformation of the visual map. *Soc Neurosci Abst* 15(1):162.

Goldberg, M. E., Colby, C. L., and Duhamel, J. R. (1990) Representation of visuomotor space in the parietal lobe of the monkey. *Cold Spring Harbor Symposia on Quantitative Biology* 55:729-739.

Goldberg, M. E., and Colby, C. L. (1992) Oculomotor control and spatial processing. *Curr Opin Neurobiol* 2:198-202.

Goodale, M. A., and Milner, A. D. (1992) Separate visual pathways for perception and action. *Trends Neurosci* 15(1):20-25.

Goodman, S. J., and Andersen, R. A. (1989) Microstimulation of a neural-network model for visually guided saccades. *J Cog Neurosci* 1(4):317-326.

Goodman, S. J., and Andersen, R. A. (1990) Algorithm programmed by a neural network model for coordinate transformation. *Proceedings of the International Joint Conference On Neural Networks*, San Diego, CA, June 1990, vol. II, 381-386.

Grüsser, O. -., Pause, M., and Schreiter, U. (1990a) Localization and responses of neurones in the parieto-insular vestibular cortex of awake monkeys (Macaca fascicularis). *J Physiol (Lond)* 430:537-557.

Grüsser, O. -., Pause, M., and Schreiter, U. (1990b) Vestibular neurones in the parieto-insular cortex of monkeys (Macaca fascicularis): visual and neck receptor responses. *J Physiol (Lond)* 430:559-583.

Gullapalli, V. (1988) *A stochastic learning algorithm for learning real-valued functions via reinforcement feedback*.COINS Technical Report 88-91, MA 01003.

Hallett, P. E., and Lightstone, A. D. (1976) Saccadic eye movements to flashed targets. *Vision Res* 16:107-114.

Hebb, D. O. (1949) *The organization of behavior*. Wiley, New York.

Hecaen, H., and De Ajuriaguerra, A. (1954) Balint's syndrome (psychic paralysis of gaze) and its minor forms. *Brain* 77:373-400.

Hecaen, H., and Albert, M. L. (1978) *Human neuropsychology*. Wiley, New York.

Hecht-Nielsen, R. (1989) Theory of the backpropagation neural network. *International Joint Conference on Neural Networks*, Washington, D.C., June 1989, vol. I, 593-605.

Heffner, H., and Masterton, R. B. (1975) Contribution of auditory cortex to sound localization in the monkey (*Macaca mulatta*). *J Neurophysiol* 38:1340-1358.

Heffner, H. E., and Heffner, R. S. (1990) Effect of bilateral auditory cortex lesions on sound localization in Japanese macaques. *J Neurophysiol* 64:915-931.

Heide, W., Zimmermann, E., and Kömpf, D. (1993) Double-step saccades in patients with frontal or parietal lesions. *Soc Neurosci Abst* 19:427.

Heilman, K. M., Pandya, D. N., and Geschwind, N. (1970) Trimodal inattention following parietal lobe ablations. *Trans. Am. Neurol. Assoc.* 95:259-261.

Heilman, K. M., Pandya, D. N., Karol, E. A., and Geschwind, N. (1971) Auditory inattention. *Arch Neurol* 24:323-325.

Heilman, K. M., and Valenstein, E. (1972) Auditory neglect in man. *Arch Neurol* 26:32-35.

Heilman, K. M., Bowers, D., Valenstein, E., and Watson, R. T. (1987) Hemispace and hemispatial neglect. In *Neurophysiological and Neuropsychological Aspects of Spatial Neglect*, M. Jeannerod, ed., pp. 115-150, Elsevier, Amsterdam.

Hikosaka, O., and Wurtz, R. H. (1983) Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. *J Neurophysiol* 49(5):1268-84.

Hikosaka, O., and Wurtz, R. H. (1985) Modification of saccadic eye movements by GABA-related substances. I. Effect of muscimol and bicuculline in monkey superior colliculus. *J Neurophysiol* 53(1):266-91.

Hikosaka, O., and Wurtz, R. H. (1989) The basal ganglia. In *The neurobiology of saccadic eye movements*, R. H. Wurtz, and M. E. Goldberg, eds., pp. 257-81, Elsevier Science Publishers BV, Amsterdam.

Hinton, G. E. (1987) *Connectionist learning procedures*.Technical Report No. CMU-CS-87-115, Pittsburgh, PA 15213.

Holmes, G. (1918) Disturbances of visual orientation. *Br J Ophthal* 2:449-468.

Holmes, G. (1938) The cerebral integration of the ocular movements. *Br. Med. J.* 2:107-112.

Husain, M., and Stein, J. (1988) Rezso Balint and his most celebrated case. *Arch Neurol* 45:89-93.

Hyvärinen, J., and Poranen, A. (1974) Function of the parietal associative area 7 as revealed from cellular discharges in alert monkeys. *Brain* 97:673-692.

Ito, M. (1984) *The cerebellum and neural control.* Raven Press, New York.

Jay, M. F., and Sparks, D. L. (1984) Auditory receptive fields in primate superior colliculus shift with changes in eye position. *Nature* 309(5966):345-347.

Jay, M. F., and Sparks, D. L. (1987a) Sensorimotor integration in the primate superior colliculus. I. Motor convergence. *J Neurophysiol* 57(1):22-34.

Jay, M. F., and Sparks, D. L. (1987b) Sensorimotor integration in the primate superior colliculus. II. Coordinates of auditory signals. *J Neurophysiol* 57(1):35-55.

Jenkins, W. M., and Masterton, R. B. (1982) Sound localization: effects of unilateral lesions in central auditory systems. *J Neurophysiol* 47:987-1016.

Jenkins, W. M., and Merzenich, M. M. (1984) Role of cat primary auditory cortex for sound-localization behavior. *J Neurophysiol* 52:819-847.

Jerger, J., Weikers, N., Sharbrough, F., and Jerger, S. (1969) Bilateral lesions of the temporal lobe: a case study. *Acta Otolaryng.* Suppl. 258:???

Jerger, J., Lovering, L., and Wertz, M. (1971) Auditory disorder following bilateral temporal lobe insult: report of a case. *Journal of Speech and Hearing Disorders* 37(4):523-535.

Judge, S. J., Richmond, B. J., and Chu, F. C. (1980) Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res* 20:535-538.

Kalaska, J. F., and Crammond, D. J. (1990) Cortical neuronal activity recorded in a delay task that dissociates location of cue stimulus and movement endpoint. *Soc Neurosci Abst* 16:423.

Kavanagh, G. L., and Kelly, J. B. (1987) Contribution of auditory cortex to sound localization by the ferret (Mustela putorius). *J Neurophysiol* 57:1746-1766.

Kelly, J. B. (1980) Effects of auditory cortical lesions on pure-tone sound localization by the albino rat. *J Neurophysiol* 44:1161-1174.

Kelso, S. R., Ganong, A. H., and Brown, T. H. (1986) Hebbian synapses in hippocampus. *Proc Natl Acad Sci USA* 83:5326-5330.

King, A. J., and Abraham (???) The integration of visual and auditory spatial information in the brain. In *Higher order sensory processing and artificial sensing*, D. M. Guthrie, ed., pp. ???, ???, ???

Klingon, G. H., and Bontecou, D. C. (1964) Auditory field localization. *Trans Amer Neurol Ass* 89:210-211.

Klingon, G. H., and Bontecou, D. C. (1966) Localization in auditory space. *Neurol* 16:879-886.

Knudsen, E. I. (1982) Auditory and visual maps of space in the optic tectum of the owl. *J Neurosci* 2(9):1177-1194.

Knudsen, E. I. (1983) Early auditory experience aligns the auditory map of space in the optic tectum of the barn owl. *Science* 222(4626):939-942.

Knudsen, E. I., and Knudsen, P. F. (1985) Vision guides the adjustment of auditory localization in young barn owls. *Science* 230(4725):545-548.

Koch, K. W., and Fuster, J. M. (1989) Unit activity in monkey parietal cortex related to haptic perception and temporary memory. *Exp Brain Res* 76:292-306.

Kosslyn, S. M. (1988) Aspects of a cognitive neuroscience of mental imagery. *Science* 240:1621-1626.

Kurylo, D. D., and Skavenski, A. A. (1991) Eye movements evoked by electrical stimulation of area PG in the monkey. *J Neurophysiol* 65(6):1243-53.

Lackner, J. R. (1973) Visual rearrangement affects auditory localization. *Neuropsychol* 11:29-32.

Leinonen L, Hyvärinen, J., and Sovijärvi, A. R. A. (1980) Functional properties of neurons in the temporo-parietal association cortex of awake monkeys. *Exp Brain Res* 39:203-215.

Lemon, R. (1988) The output map of the primate motor cortex. *Trends Neurosci* 11:501-506.

Linsker, R. (1989) *How to generate ordered maps by maximizing the mutual information between input and output signals.* IBM Technical Report No. 65530, Yorktown Heights, NY.

Lippmann, R. P. (1987) An introduction to computing with neural nets. *IEEE ASSP magazine* (4):4-22.

Lynch, J. C., Mountcastle, V. B., Talbot, W. H., and Yin, T. C. T. (1977) Parietal lobe mechanisms for directed visual attention. *J Neurophysiol* 40:362-389.

Lynch, J. C. (1980) The functional organization of posterior parietal association cortex. *Behav Brain Sci* 3:485-499.

Lynch, J. C. (1980) The role of the parieto-occipital association cortex in oculomotor control. *Exp Brain Res* 41:A32.

Lynch, J. C., and McLaren, J. W. (1983) Optokinetic nystagmus deficits following parieto-occipital cortex lesions in monkeys. *Exp Brain Res* 49:125-130.

Lynch, J. C., Graybiel, A. M., and Lobeck, L. J. (1985) The differential projection of two cytoarchitectonic subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus. *J Comp Neurol* 253:241-254.

Lynch, J. C., and McLaren, J. W. (1989) Deficits of visual attention and saccadic eye movements after lesions of parietoo pital cortex in monkeys. *J Neurophysiol* 61(1):74-90.

Lynch, J. C. (1992) Saccade initiation and latency deficits after combined lesions of the frontal and posterior eye fields in monkeys. *J Neurophysiol* 68(5):1913-1916.

Mays, L. E., and Sparks, D. L. (1980) Dissociation of visual and saccade-related responses in superior colliculus neurons. *J Neurophysiol* 43(1):207-232.

Mays, L. E., and Sparks, D. L. (1980) Saccades are spatially, not retinocentrically, coded. *Science* 208:1163-1165.

Mazzoni, P., and Andersen, R. A. (in press) Coordinate transformations for gaze coding in the posterior parietal cortex. In *The Handbook of Brain Theory and Neural Networks*, M. A. Arbib, ed., MIT Press, Cambridge, MA.

Mazzoni, P., Andersen, R. A., and Jordan, M. I. (1990) A_{R-P} learning applied to a network model of cortical area 7a. *Proceedings of the International Joint Conference On Neural Networks*, San Diego, CA, June 1990, vol. II, 373-379.

Mazzoni, P., Andersen, R. A., and Jordan, M. I. (1991a) A more biologically plausible learning rule than backpropagation applied to a network model of cortical area 7a. *Cerebral Cortex* 1(4):293-307.

Mazzoni, P., Andersen, R. A., and Jordan, M. I. (1991b) A more biologically plausible learning rule for neural networks. *Proc Natl Acad Sci USA* 88:4433-4437.

Mazzoni, P., Bracewell, R. M., Barash, S., and Andersen, R. A. (1992) Intended movement activity of area LIP neurons in delayed double saccade tasks. *Soc Neurosci Abst* 18:148.

Mazzoni, P., and Andersen, R. A. (1993) The role of neural networks in the study of the posterior parietal cortex. *Proceedings of the International Conference On Neural Networks*, San Francisco, CA, March 1993, vol. III, 1321-1326.

Mazzoni, P., Bracewell, R. M., Barash, S., and Andersen, R. A. (1993) Spatially tuned auditory responses in area LIP of macaques

performing memory saccades to acoustic targets . *Soc Neurosci Abst* 19:26.

McClelland, J. L., and Rumelhart, D. E. (1988) *Explorations in parallel distributed processing*. MIT Press, Cambridge, MA.

Middlebrooks, J. C., and Green, D. M. (1991) Sound localization by human listeners. *Annu Rev Psychol* 42:135-159.

Minsky, M. L., and Papert, S. (1969) *Perceptrons*. MIT Press, Cambridge.

Mishkin, M., Ungerleider, L. G., and Macko, K. A. (1983) Object vision and spatial vision: two cortical pathways. *Trends Neurosci* ???[***](10):414-417.

Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H., and Acuna, C. (1975) Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *J Neurophysiol* 38(4):871-908.

Neff, W. D., Fisher, J. F., Diamond, I. T., and Yela, M. (1956) Role of auditory cortex in discrimination requiring localization of sound in space. *J Neurophysiol* 19:500-512.

Palmer, A. R., and King, A. J. (1982) The representation of auditory space in the mammalian superior colliculus. *Nature* 299:248-249.

Pandya, D. N., and Kuypers H G J M (1969) Cortico-cortical connections of the rhesus monkey. *Brain Res* 13:13-36.

Pandya, D. N., and Sanides, F. (1973) Architectonic parcellation of the temporal operculum in rhesus monkey and its projection pattern. *Z Entwicklungsgesch* 139:127-161.

Pandya, D. N., and Seltzer, B. (1982) Intrinsic connections and architectonics of posterior parietal cortex in the rhesus monkey. *J Comp Neurol* 204:196-210.

Pandya, D. N., and Yeterian, E. H. (1985) Architecture and connections of cortical association areas. In *Cerebral Cortex Vol. 4 Association and auditory cortices*, A. Peters, and E. G. Jones, eds., Plenum, NY.

Parker, D. B. (1985) *Learning logic.*Technical Report TR-47, Cambridge, MA 02139.

Phillips, D. P. (1985) Progress in neurophysiology of sound localization. *Annu Rev Psychol* 36:245-274.

Pick, H. L., Warren, D. H., and Hay, J. C. (1969) Sensory conflict in judgments of spatial direction. *Percept Psychophys* 6(4):203-205.

Pierrot-Deseilligny, C., Rivaud, S., Gaymard, B., and Agid, Y. (1991) Cortical control of memory saccades in man. *Exp Brain Res* 83(3):607-617.

Pinek, B., Duhamel, J. R., Cavé, C., and Brouchon, M. (1989) Audio-spatial deficits in humans: differential effects associated with left versus right hemisphere parietal damage. *Cortex* 25:175-186.

Platt, B. B., and Warren, D. H. (1972) Auditory localization: the importance of eye movements and a textured visual environment. *Percept Psychophys* 12(2B):245-248.

Posner, M. I., Walker, J. A., Friedrich, F. J., and Rafal, R. D. (1984) Effects of parietal lobe injury on covert orienting of visual attention. *J Neurosci* 4:1863-1874.

Quintana, J., and Fuster, J. M. (1993) Spatial and temporal factors in the role of prefrontal and parietal cortex in visuomotor integration. *Cerebral Cortex* 3:122-132.

Radeau, M. (1973) The locus of adaptation to auditory-visual conflict. *Perception* 2:327-332.

Rayleigh, L. (1876) Our perception of the direction of a source of sound. *Nature* 14:32-33.

Rayleigh, L. (1907) On our perception of sound direction. *Philos Mag* 13:214-232.

Richardson, R. T., Mitchell, S. J., Baker, F. H., and DeLong, M. R. (1988) Responses of nucleus basalis of Meynert neurons in behaving monkeys. In *Cellular mechanisms of conditioning and behavioral plasticity*, C. D. Woody, D. L. Alkon, and J. L. McGaugh, eds., pp. 161-173, Plenum, New York.

Robinson, D. A. (1963) A method for measuring eye movement using a scleral search coil in a magnetic field. *IEEE Trans. Biomed. Eng.* 10:137-145.

Robinson, D. L., Goldberg, M. E., and Stanton, G. B. (1978) Parietal association cortex in the primate: sensory mechanisms and behavioral modulation. *J Neurophysiol* 41:910-932.

Ruff, R. M., Hersh, N. A., and Pribram, K. H. (1981) Auditory spatial deficits in the personal and extrapersonal frames of reference due to cortical lesions. *Neuropsychol* 19:435-443.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986) Learning internal representations by error propagation. In *Parallel distributed processing: explorations in the microstructure of cognition*, D. E. Rumelhart, and J. L. McClelland, eds., vol. I, pp. 318-362, MIT Press, Cambridge, MA.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986b) Learning representations by back-propagating errors. *Nature* 323:533-536.

Sakata, H., Takaoka, Y., Kawarasaki, A., and Shibutani, H. (1973) Somatosensory properties of neurons in the superior parietal cortex (area 5) of the rhesus monkey. *Brain Res* 64:85-102.

Sanchez-Longo, L. P., Forster, F. M., and Auth, T. L. (1957) A clinical test for sound localization and its applications. *Neurol* 7:655-663.

Sanchez-Longo, L. P., and Forster, F. M. (1958) Clinical significance of impairment of sound localization. *Neurol* 8:119-125.

Sawaguchi, T., and Goldman-Rakic, P. S. (1991) D1 dopamine receptors in prefrontal cortex: involvement in working memory. *Science* 251:947-950.

Schiller, P. H., True, S. D., and Conway, J. L. (1980) Deficits in eye movements following frontal eye-field and superior colliculus ablations. *J Neurophysiol* 44:1175-89.

Schiller, P. H. (1984) The superior colliculus and visual function. In *The Handbook of Physiology. Section 1: The Nervous System, Vol. III. Sensory Processes, Part 1*, J. M. Brookhart, V. B. Mountcastle, I. Darian-Smith, and S. T. Geiger, eds., pp. 457-505, American Physiological Society, Bethesda, MD.

Seal, J., Gross, C., and Bioulac, B. (1982) Activity of neurons in area 5 during a simple arm movement in monkeys before and after deafferentiation of the trained limb. *Brain Res* 250:229-243.

Seal, J., Gross, C., Doudet, D., and Bioulac, B. (1983) Instruction-related changes of neuronal activity in area 5 during simple forearm movement in the monkey. *Neurosci Lett* 36:145-150.

Sejnowski, T. J. (1981) Skeleton filters in the brain. In *Parallel models of associative memory*, G. E. Hinton, and J. A. Anderson, eds., pp. 189-212, Lawrence Erlbaum Associates, Hillsdale, NJ.

Sejnowski, T. J., Rosenberg, C. R. (1986) *NETtalk: a parallel network that learns to read aloud.*Technical Report JHU/EECS-86/01, The Johns Hopkins University Electrical Engineering and Computer Science.

Sejnowski, T. J., Chattarji, S., and Stanton, P. K. (1989) Introduction of synaptic plasticity by Hebbian covariance in the hippocampus. In *The computing neuron*, R. M. Durbin, Miall R C, and G. J. Mitchison, eds., pp. 105-124, Addison-Wesley, New York.

Shankweiler, D. P. (1961) Performance of brain-damaged patients on two tests of sound localization. *J Comp Physiol Psychol* 54(4):375-381.

Shea, P. M., and Lin, V. (1989) Detection of explosives in checked airline baggage using an artificial neural system. *International*

Joint Conference On Neural Networks, Washington, D.C., June 1989, vol. II, 31-34.

Shibutani, H., Sakata, H., and Hyvarinen, J. (1984) Saccade and blinking evoked by microstimulation of the posterior parietal association cortex of the monkey. *Exp Brain Res* 55(1):1-8.

Snyder, L. H., Brotchie, P., and Andersen, R. A. (1993) World-centered encoding of location in posterior parietal cortex of monkey. *Soc Neurosci Abst* 19:770.

Sparks, R., Goodglass, H., and Nickel, B. (1970) Ipsilateral versus contralateral extinction in dichotic listening resulting from hemisphere lesions. *Cortex* 6:249-260.

Sparks, D. L. (1986) The neural translation of sensory signals into commands for control of saccadic eye movements: the role of primate superior colliculus. *Psychol Rev* 66(1):118-171.

Sparks, D. L., and Hartwich-Young, R. (1989) The deep layers of the superior colliculus. In *The Neurobiology of Saccadic Eye Movements*, R. H. Wurtz, and M. E. Goldberg, eds., pp. 213-255, Elsevier Science Publishers BV, Amsterdam.

Stanford, T. R., and Sparks, D. L. (1994) Systematic error for saccades to remembered targets: evidence for a dissociation between saccade metrics and activity in the superior colliculus. *Vision Res* 34(1):93-106.

Stanton, P. K., and Sejnowski, T. J. (1989) Associative long-term depression in the hippocampus induced by hebbian covariance. *Nature* 339:215-218.

Stein, J. F. (1976) The effect of cooling parietal lobe areas 5 and 7 upon voluntary movement in awake rhesus monkeys. *J Physiol (Lond)* 258:62-63P.

Stein, B. E., Meredith, M. A., Huneycutt, W. S., and McDade, L. (1989) Behavioral indices of multisensory integration: orientation to visual cues is affected by auditory stimuli. *J Cog Neurosci* 1(1):12-24.

Syka, J., and Masterton, R. B. (eds.) (1988) *Auditory pathways: structure and function*, Plenum, New York.

Thier, P., and Andersen, R. A. (1991) Electrical microstimulation delineates 3 distinct eye-movement related areas in the posterior parietal cortex of the rhesus monkey. *Soc Neurosci Abst* 17:1281.

Tolhurst, D. J., Movshon, J. A., and Dean, A. F. (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23(8):775-785.

Tolhurst, D. J. (1989) The amount of information transmitted about contrast by neurones in the cat's visual cortex. *Visu Neurosci* 2:409-413.

Ungerleider, L. G., and Mishkin, M. (1982) Two cortical visual systems. In *The Analysis of Visual Behavior*, D. A. Ingle, M. A. Goodale, and R. J. W. Mansfield, eds., MIT, Cambridge, MA.

Vogels, R., Spileers, W., and Orban, G. A. (1989) The response variability of striate cortical neurons in the behaving monkey. *Exp Brain Res* 77:432-436.

Walsh, E. G. (1957) An investigation of sound localization in patients with neurological abnormalities. *Brain* 80:222-250.

Warren, D. H., and Pick, H. L. (1970) Intermodality relations in localization in blind and sighted people. *Percept Psychophys* 8(6):430-432.

Weinrich, M., Wise, S. P., and Mauritz, K. (1984) A neurophysiological analysis of the premotor cortex in the monkey. *Brain* 107:385-414.

Welch, R. B., and Warren, D. H. (1986) Intersensory interactions. In *Handbook of perception and human performance*, K. R. Boff, L. Kaufman, and J. P. Thomas, eds., pp. 25.1-25.36, Springer-Verlag, Berlin

Werbos, P. J. (1974). *Beyond regression: new tools for prediction and analysis in the behavioral sciences.* Unpublished doctoral dissertation, Harvard University, Cambridge, MA.

White, J. M., Sparks, D. L., and Stanford, T. R. (1994) Saccades to remembered target locations: and analysis of systematic and variable errors. *Vision Res* 34(1):79-92.

Widrow, B., and Hoff, M. E. (1960) Adaptive switching circuits. *Institute of Radio Engineers, Western Electronic Show and Convention, Convention Record,* vol. Part 4, 96-104.

Williams, R. J. (1986) *Reinforcement learning in connectionist networks: a mathematical analysis.*Technical Report ICS 8605, La Jolla, CA 92093.

Williams, R. J. (1987) A class of gradient-estimating algorithms for reinforcement learning in neural networks. *Proceedings of IEEE First Annual International Conference on Neural Networks*, San Diego, CA, June 1987, vol. II, 601-608.

Wise, S. P., Weinrich, M., and Mauritz, K. (1983) Motor aspects of cue-related neuronal activity in premotor cortex of the rhesus monkey. *Brain Res* 260:301-305.

Wise, S. P. (1985) The primate premotor cortex: past, present and preparatory. *Annu Rev Neurosci* 8(1-19):

Wise, S. P., and Mauritz, K. (1985) Set related neuronal activity in the premotor cortex of rhesus monkeys: effects of changes in motor set. *Proc R Soc Lond Ser B* B223:331-354.

Wolbarsht, M. L., MacNichol, F. F. J., and Wagner, H. G. (1960) Glass insulated platinum microelectrode. *Science* 132:1309-1310.

Wurtz, R. H., and Goldberg, M. E. (1989) *The neurobiology of saccadic eye movements.* Elsevier Science Publishers BV, Amsterdam.

Yost, W. A. , and Gourevitch, G. (eds.) (1987) *Directional hearing*, Springer-Verlag, Berlin.

Zipser, D., and Andersen, R. A. (1988) A backpropagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679-684.

Zipser, D., and Rumelhart, D. E. (1990) Neurobiological significance of new learning models. In *Computational neuroscience*, E. L. Schwartz, ed., pp. 192-200, MIT Press, Cambridge, MA.