

**Responses to auditory stimuli
in macaque lateral intraparietal area**

Thesis by
Jennifer F. Linden

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

1999
(Submitted May 14, 1999)

© 1999

Jennifer F. Linden

All Rights Reserved

The end of our exploring
will be to arrive at where we started,
and to know the place for the first time.

T. S. Eliot

This thesis is dedicated to my parents,
Betty and Ted Linden.

Acknowledgements

My years as a graduate student have been very enjoyable, thanks to many wonderful people at Caltech.

First and foremost, I thank my advisor Richard Andersen, for his support, encouragement, and guidance over the past five years. I feel lucky to have had the opportunity to participate in the vibrant and exciting intellectual life of his laboratory.

I also thank the members of my candidacy and thesis committees — Gilles Laurent, Mark Konishi, Pietro Perona, Christof Koch, and Scott Fraser — for their help and advice throughout my graduate career at Caltech. I am especially indebted to Gilles, Christof, and Scott, for their guidance during the rotations I did in their laboratories during my first two years at Caltech. These laboratory rotations contributed greatly to my scientific education, by introducing me to a wide range of research questions and experimental techniques.

Special acknowledgements and extra thanks are due to three people in the Andersen lab who helped make this thesis possible: Alexander Grunewald, Maneesh Sahani, and Larry Snyder. Alexander collaborated with me on data collection for this thesis, and on many of the analyses presented in Chapters 2 and 3. Maneesh collaborated with me on development and application of the spike train analysis algorithms discussed in Chapter 4. Maneesh also wrote the data collection software package used to run many of the experiments in the laboratory, including those described in this thesis. Larry, while not directly involved in this research, laid the foundations for it by teaching me how to do extracellular recordings in awake behaving monkeys, and by providing me with invaluable advice on behavioral training and experimental design.

I also thank Betty Gillikan and Deborah Ward for technical help with animal care and surgical procedures; Cierina Reyes, Sylvie Gertamenian, and Laura Rodriguez

for administrative assistance; Viktor Shcherbatyuk for computer support; and Janet Baer for veterinary advice.

I am grateful to the Howard Hughes Medical Institute for providing me with the Predoctoral Fellowship which funded the first five years of my graduate work. Support for the last two years of my research came from the National Institutes of Health.

In addition to the people mentioned above, I thank all the present and past members of the Andersen lab, for stimulating discussions, constructive criticism, practical advice, technical assistance, and comic relief. I also thank all of my contemporaries in the Sloan Center for Theoretical Neurobiology, the Computation and Neural Systems program, and other graduate programs at Caltech, for teaching me about many areas of science beyond the scope of my own research.

Finally, I thank my family and friends, for providing the moral support which has enabled me to reach this point. I am especially grateful to my wonderful parents, Betty and Ted Linden, whose faith in me has been my greatest source of strength. I am also deeply thankful for the presence in my life of my beloved aunts, Patricia Plante and Mary-Ann Myrant; my wonderful brother and sister-in-law, Greg and Corina Linden; and my dearest friend Maneesh.

Abstract

The lateral intraparietal area (LIP) of macaque posterior parietal cortex participates in the sensorimotor transformations underlying visually guided eye movements. Area LIP has long been considered unresponsive to auditory stimulation. However, recent studies have shown that neurons in LIP respond to auditory stimuli during an auditory-saccade task, suggesting possible involvement of this area in auditory-to-oculomotor as well as visual-to-oculomotor processing. This dissertation describes investigations which clarify the role of area LIP in auditory-to-oculomotor processing.

Extracellular recordings were obtained from a total of 332 LIP neurons in two macaque monkeys, while the animals performed fixation and saccade tasks involving auditory and visual stimuli. No auditory activity was observed in area LIP before animals were trained to make saccades to auditory stimuli, but responses to auditory stimuli did emerge after auditory-saccade training. Auditory responses in area LIP after auditory-saccade training were significantly stronger in the context of an auditory-saccade task than in the context of a fixation task. Compared to visual responses, auditory responses were also significantly more predictive of movement-related activity in the saccade task. Moreover, while visual responses often had a fast transient component, responses to auditory stimuli in area LIP tended to be gradual in onset and relatively prolonged in duration.

Overall, the analyses demonstrate that responses to auditory stimuli in area LIP are dependent on auditory-saccade training, modulated by behavioral context, and characterized by slow-onset, sustained response profiles. These findings suggest that responses to auditory stimuli are best interpreted as supramodal (cognitive or motor) responses, rather than as modality-specific sensory responses. Auditory responses in area LIP seem to reflect the significance of auditory stimuli as potential targets for eye movements, and may differ from most visual responses in the extent to which they are abstracted from the sensory parameters of the stimulus.

Contents

Acknowledgements	v
Abstract	vii
Organization of the Thesis	xiv
1 Introduction	1
1.1 Posterior parietal cortex	2
1.2 Lateral intraparietal area (LIP)	6
1.3 Sound localization	11
1.3.1 Barn owls	11
1.3.2 Primates	14
1.4 Saccadic eye movements	16
1.5 Sound localization and eye movements	18
1.6 Responses to auditory stimuli in area LIP	20
2 Training Effects	22
2.1 Introduction	23
2.2 Methods	25
2.2.1 Animals and animal care	25
2.2.2 Surgical procedures	26
2.2.3 Experimental setup	26
2.2.4 Behavioral paradigms	28
2.2.5 Recording procedures	32
2.2.6 Recording strategies	32
2.2.7 Database	33
2.2.8 Data analysis	34

2.2.9	Histology	36
2.3	Results	37
2.3.1	Pre-training	37
2.3.2	Post-training	44
2.3.3	Comparison between pre- and post-training	44
2.3.4	Control for search/selection bias	46
2.3.5	Control for different arrays	47
2.3.6	Control for number of locations/samples	48
2.3.7	Control for post-reward eye movements	49
2.3.8	Control for penetration locations	51
2.4	Discussion	51
2.4.1	Effects of auditory-saccade training	51
2.4.2	Possible mechanisms	54
2.4.3	Relation to previous work	55
2.4.4	Interpretations	56
3	Behavioral Modulation	60
3.1	Introduction	61
3.2	Methods	62
3.2.1	Animals, animal care, and surgical procedures	62
3.2.2	Experimental setup	62
3.2.3	Behavioral paradigms	63
3.2.4	Recording procedures and strategy	66
3.2.5	Database	67
3.2.6	Data analysis	68
3.2.7	Histology	69
3.3	Results	69
3.3.1	Behavioral modulation: stimulus period	69
3.3.2	Behavioral modulation: later periods	77
3.3.3	Auditory-visual correlation: stimulus period	82

3.3.4	Auditory-visual correlation: later periods	85
3.3.5	Correlation between periods	85
3.3.6	Control for response measure	88
3.3.7	Control for block order	90
3.4	Discussion	91
3.4.1	Behavioral modulation of auditory responses	92
3.4.2	No behavioral modulation of visual responses?	92
3.4.3	Behavioral modulation of delay activity	93
3.4.4	Association between auditory and visual responses	95
3.4.5	Link between auditory and delay/saccade activity	96
3.4.6	Experimental considerations	96
3.4.7	Interpretations	99
4	Temporal Features	102
4.1	Introduction	103
4.1.1	Subintervals	103
4.1.2	Spike density functions	104
4.1.3	Distance metrics	106
4.1.4	Hidden Markov models	107
4.1.5	Randomly scaled inhomogeneous Poisson model	108
4.2	Methods	110
4.2.1	Database and data analysis	110
4.2.2	Probabilistic generative model	111
4.2.3	Mixture models	116
4.2.4	Model selection	121
4.3	Results	122
4.3.1	Population histograms	122
4.3.2	Variability across trials	125
4.3.3	Goodness-of-fit	132
4.3.4	Response latencies	137

4.3.5	Temporal features	143
4.4	Discussion	158
4.4.1	Temporal features of auditory and visual responses	158
4.4.2	Randomly scaled inhomogeneous Poisson model	161
4.4.3	Possible improvements	164
5	Conclusions	167
5.1	Responses to auditory stimuli in area LIP	168
5.2	Implications	168
5.2.1	Role of LIP in auditory-to-oculomotor processing	169
5.2.2	Function of area LIP	170
5.2.3	Auditory-to-oculomotor transformation in primates	171
5.3	Future directions	172

List of Figures

2.1	Pre- and post-training fixation tasks and stimulus arrays.	29
2.2	Electrode track through the intraparietal sulcus	38
2.3	Example of neural activity before auditory-saccade training.	39
2.4	Percentages of cells with spatially tuned auditory or visual responses.	41
2.5	Most significant auditory response in pre-training database.	42
2.6	Histograms of stimulus location information.	43
2.7	Example of neural activity after auditory-saccade training.	45
2.8	Example of post-reward eye position distributions.	50
2.9	Penetration maps for fixation task.	52
3.1	Memory-saccade task, fixation task, and stimulus array.	64
3.2	Example of neural activity during the two auditory tasks.	71
3.3	Example of neural activity during the two visual tasks.	72
3.4	Behavioral modulation in the stimulus and pre-stimulus periods.	73
3.5	Behavioral modulation of weakly tuned visual responses.	76
3.6	Example of neural activity during the memory-saccade task.	78
3.7	Behavioral modulation in the delay and saccade/hold periods.	80
3.8	Auditory-visual correlations in the stimulus and pre-stimulus periods.	83
3.9	Penetration maps for the memory-saccade task.	84
3.10	Auditory-visual correlations in the delay and saccade/hold periods.	86
3.11	Correlations across response periods.	89
4.1	The probabilistic generative model.	112
4.2	A simulated mixture model.	118
4.3	Population histograms.	124
4.4	Relationship between spike count variance and spike count mean.	127
4.5	Comparison of multiplicative and additive scaling hypotheses.	131

4.6	Fits of the generative model to real data.	135
4.7	Simulated data with re-fits.	136
4.8	Model fit percentile histograms.	138
4.9	Estimated response latencies for four different cells.	140
4.10	Response latency histograms.	141
4.11	Simulated database with four temporal features.	145
4.12	Optimal mixture model for simulated database.	146
4.13	Optimal mixture model for auditory memory-saccade task.	148
4.13	Optimal mixture model for auditory memory-saccade task, continued.	149
4.14	Optimal mixture model for visual memory-saccade task.	150
4.14	Optimal mixture model for visual memory-saccade task, continued. .	151
4.14	Optimal mixture model for visual memory-saccade task, continued. .	152
4.15	Optimal mixture model for auditory fixation task.	154
4.16	Optimal mixture model for visual fixation task.	155
4.16	Optimal mixture model for visual fixation task, continued.	156
4.16	Optimal mixture model for visual fixation task, continued.	157

Organization of the Thesis

This dissertation is organized into five chapters, of which three present original results¹ concerning responses to auditory stimuli in macaque parietal cortex. A brief chapter summary is provided on the first page of each chapter. The three central chapters (Training Effects, Behavioral Modulation, and Temporal Features) address different questions about auditory responses in the lateral intraparietal area, and therefore contain chapter-specific introductions and discussions. The remaining two chapters (Introduction and Conclusions) provide a more general overview of the context of the work and the possible implications of the results.

Chapter 1 (Introduction) begins with a brief review of the anatomy and physiology of the posterior parietal cortex, with particular focus on the lateral intraparietal area. Debates over the function of the lateral intraparietal area in visual and oculomotor processing are summarized, and recent findings suggesting a possible role for area LIP in auditory processing are presented. The chapter continues with reviews of sound localization and saccadic eye movement systems in primates, and closes with some speculation on possible interactions between the auditory and oculomotor systems.

Chapter 2 (Training Effects) presents the first neurophysiological study in this thesis: an investigation of the effects of training on auditory responses in area LIP. This chapter includes a summary of training effects observed in other areas, and a discussion of the possible implications of training-induced changes in area LIP.

Chapter 3 (Behavioral Modulation) explores the dependence of auditory responses in area LIP on behavioral context. This chapter also addresses the relationship between auditory and visual responses, and the link between auditory responses and movement-related activity. A brief review of movement-related auditory responses found in other parts of the brain is provided in the chapter discussion.

¹Please see Acknowledgements for a summary of the contributions of other researchers to the work which appears in this thesis.

Chapter 4 (Temporal Features) presents a detailed comparison of auditory and visual response profiles in area LIP. The chapter begins with a review of existing methods for analyzing the temporal features of spike trains, and then outlines a novel and more principled approach to the problem. A doubly stochastic Poisson model for cortical spike trains is described in the Methods section, and algorithms for smoothing and clustering spike trains are derived from this model. These algorithms are then applied to data recorded from area LIP, to identify differences in the temporal features of responses to auditory and visual stimuli.

Finally, Chapter 5 (Conclusions) pulls together the results of Chapters 2, 3, and 4, to propose a new interpretation of responses to auditory stimuli in area LIP. This chapter also offers some speculations on the role of LIP in auditory-to-oculomotor processing, the overall function of area LIP, and the process of auditory-to-oculomotor transformation in primates. The chapter and thesis end with suggestions for future investigations.

Chapter 1 Introduction

The primate posterior parietal cortex seems to be specialized for implementing sensorimotor transformations — that is, for transforming sensory information into motor commands. In particular, the lateral intraparietal area (LIP) within the inferior parietal lobule plays a major role in visual-to-oculomotor transformations. Anatomically, area LIP is well situated to participate in planning visually guided eye movements; the area receives strong inputs from extrastriate visual areas, and projects to cortical and subcortical oculomotor control centers. The physiology of LIP further indicates that this area integrates visual information with eye position signals to direct saccades to visual targets. Recent physiological evidence suggests that area LIP may be involved in auditory-to-oculomotor as well as visual-to-oculomotor transformations. Although early studies of the region reported no responses to auditory stimulation, more recent investigations have found that LIP neurons respond to auditory stimuli in the context of a saccade task. This finding raises the possibility that area LIP might participate in linking the sound localization system and the saccadic eye movement system in the primate brain. Previous behavioral and neurophysiological results suggest that extensive interconnections between these two systems exist; however, it is not known if area LIP is actually involved in auditory-to-oculomotor processing. This dissertation describes investigations of responses to auditory stimuli in area LIP. The experiments provide insight into the role of area LIP in directing saccades to auditory targets, the function of area LIP overall, and the process of auditory-to-oculomotor transformation in primates.

1.1 Posterior parietal cortex

Most behaviors require some kind of sensorimotor transformation; information is acquired by a sensory modality, and based on that information motor acts are executed. A complete understanding of how the brain controls behavior can only be achieved once the process of sensorimotor transformation is understood. For goal-directed movement in particular, this process presumably occurs in multiple stages. At the sensory end, information must be acquired through sensory transduction, and then processed to extract the location of the stimulus. At the motor end, movement output must be generated through the coordinated activation of the muscles. In between these two extremes, several other steps necessary for sensorimotor transformation may be identified. For example, attention is directed toward the stimulus; the stimulus is recognized as a potential target for movement; a decision to move is made; and the location of the target is transformed from sensory to motor coordinates.

The primate posterior parietal cortex has long been recognized as critical for sensorimotor transformations, based on the consequences of posterior parietal damage in humans. Patients with lesions in posterior parietal cortex, in particular in the inferior parietal lobule, tend to have complex deficits in spatial perception, spatial attention, and sensorimotor integration (for comprehensive reviews, see Critchley 1966; Hyvärinen 1982a; Hyvärinen 1982b; Andersen 1987; Andersen and Gnadt 1989; Stein 1989). These patients have no low-level sensory or motor problems — they are not blind, deaf, or paralyzed — yet they have difficulty using sensory information to perceive spatial relationships and to direct appropriate movements in space. Common consequences of injury to the inferior parietal lobule include:

- spatial memory deficits (impaired route-finding ability, and difficulty recalling topographic relationships);
- constructional apraxia (inability to reproduce spatial relationships between objects);
- unilateral neglect (lack of attention to the region of space contralateral to the

lesion);

- psychic paralysis of gaze (difficulty in disengaging fixation and in making voluntary eye movements);
- optic ataxia/apraxia and spatial disorientation (inability to locate objects in space); and
- extinction or simultanagnosia (failure to perceive a stimulus in the region of space contralateral to the lesion when a second stimulus is presented simultaneously on the ipsilateral side).

The last three symptoms are the features of a clinical condition known as Balint's syndrome, after the German neurologist who first described it in a patient with bilateral posterior parietal cortex damage (see Hécaen and de Ajuriaguerra 1954 for a summary in English of the original German paper, Balint 1909).

Balint's syndrome and many of the other deficits listed above are typically characterized with respect to visual stimuli, but parietal lesions may also be associated with deficits in the auditory domain. In particular, patients with damage to the inferior parietal lobule tend to mislocalize contralateral sounds into the region of space ipsilateral to the lesion (e.g., Bisiach et al. 1984; Pinek et al. 1989). This phenomenon seems to reflect a compression or distortion of perceived auditory space toward the ipsilateral side (Vallar et al. 1995). These auditory deficits, along with many of the visuospatial and oculomotor abnormalities described above, have been reproduced in monkeys with parietal lesions (Heilman et al. 1971; Lynch and McLaren 1989; Lynch 1992; Hyvärinen 1982b).

The anatomy of monkey posterior parietal cortex resembles that of the human (Hyvärinen 1982a; Andersen 1987). As a whole, posterior parietal cortex receives inputs from multiple sensory modalities, and is anatomically linked to frontal, temporal, and limbic association areas. Within macaque posterior parietal cortex, the inferior parietal lobule (area 7 of Brodmann 1905) extends from the lateral bank of the intraparietal sulcus across the caudal part of the superior temporal sulcus, and up

to the medial wall of the cerebral hemisphere (Andersen 1987). This region has extensive connections with several extrastriate visual areas, and also with somatosensory regions, the superior temporal polysensory area, prefrontal cortex, the hippocampal formation, and cingulate cortex (Andersen 1987; Hyvärinen 1982b; Seltzer and Pandya 1984). Thus the inferior parietal lobule forms a node in a network of connections between cortical areas involved in vision, somatosensation, polysensory integration, memory, motivation, and movement planning. The inferior parietal lobule also receives inputs from the pulvinar, a thalamic nucleus thought to mediate attention (Asanuma et al. 1985; Andersen 1987), and projects to the basal ganglia and other subcortical areas involved in movement control (Hyvärinen 1982a).

The inferior parietal lobule seems to be parcellated into several distinct functional areas. On cytoarchitectural grounds, Vogt and Vogt (1919) divided Brodmann’s area 7 into a caudal-medial area 7a and a rostral-lateral area 7b; these regions correspond to areas PG and PF of Von Bonin and Bailey (1947). Finer subdivisions have since been proposed on the basis of myelination differences, anatomical connections, and physiological features. The list below summarizes the anatomical and physiological characteristics of the major subregions in the inferior parietal lobule.

- **Area 7b.** Located on the rostral-lateral part of the gyrus between the intraparietal and superior temporal sulci, area 7b receives major input from somatosensory cortical areas (such as the insular cortex, area SII, and area 5) and projects primarily to premotor areas in frontal cortex (Andersen 1987; Andersen et al. 1990). Neurons in area 7b are responsive to somatosensory stimulation (Hyvärinen 1982a), and are active during reaching and hand manipulation (Mountcastle et al. 1975; Andersen 1987). This area is therefore presumed to be involved in controlling limb movements. Recently, an area which may be more specifically involved in fine hand manipulations has been identified in the cortex adjacent to area 7b along the posterior bank of the intraparietal sulcus (Taira et al. 1990); this region has been designated AIP, the anterior intraparietal area (Sakata et al. 1995).

- **Area MST.** The medial superior temporal area (MST) lies in the anterior bank of the superior temporal sulcus, and receives direct projections from several extrastriate visual areas, especially the motion-sensitive middle temporal area MT (Maunsell and Van Essen 1983). Within the posterior parietal cortex, area MST projects primarily to areas 7a and LIP (Andersen 1987). Neurons in area MST respond to smooth-pursuit eye movements and also to complex full-field visual motion such as expansion, compression, and rotation; this area is therefore thought to play an important role in direction-of-heading computations (Bradley et al. 1996; Andersen 1997).
- **Area VIP.** The ventral intraparietal area (VIP) sits in the fundus of the intraparietal sulcus, ventral to area LIP. This area is heavily interconnected with motion processing area MT (Maunsell and Van Essen 1983), and is also linked to somatosensory cortex (Seltzer and Pandya 1986). Neurons in VIP respond to localized and full-field motion (Colby et al. 1993; Schaafsma and Duysens 1996), especially motion of objects toward the face, and often have somatosensory receptive fields on the face (Duhamel et al. 1998). Cells in this region have also been reported to have spatially invariant (non-retinotopic) visual receptive fields (Duhamel et al. 1997). Recent studies have proposed that this area is involved in representation of objects within near extrapersonal space (Duhamel et al. 1998).
- **Area 7a.** A smaller region than that originally defined by Vogt and Vogt (1919) on cytoarchitectural criteria, area 7a covers the caudal-medial portion of the gyrus between the intraparietal and superior temporal sulci. Area 7a has extensive connections with extrastriate visual areas in the occipital and parietal lobes; with association areas in the frontal and temporal lobes; with limbic regions in the cingulate gyrus; and with regions of the superior temporal sulcus, including area MST, the superior temporal polysensory area, and the inferotemporal cortex (Andersen 1987; Andersen et al. 1990). Area 7a projects strongly to dorsolateral prefrontal cortex, but only weakly to the frontal eye

fields (Andersen 1987; Andersen et al. 1990). Neurons in area 7a have large, bilateral visual receptive fields, with visual responses that are modulated by eye position (Andersen and Mountcastle 1983; Andersen et al. 1985; Andersen et al. 1987) and often associated with saccade-related activity (Andersen et al. 1990; Barash et al. 1991a). Area 7a has been proposed as a high-level link between the dorsal and ventral streams in the visual system (Ungerleider and Mishkin 1982; Andersen 1987; Andersen et al. 1990), and along with area LIP, is thought to be involved in coordinate transformations for eye movements (Andersen et al. 1993).

- **Area LIP.** The lateral intraparietal area lies in the lateral bank of the intraparietal sulcus, and is the focus of this dissertation. The anatomy and physiology of this area are described in detail below.

1.2 Lateral intraparietal area (LIP)

The lateral intraparietal area is located at the top of the dorsal visual stream, the “where” pathway in vision (Ungerleider and Mishkin 1982). Based its cortico-cortical connections, area LIP seems ideally situated for involvement in visual-to-oculomotor processing. Area LIP was originally distinguished from area 7a by its strong projection to prefrontal oculomotor area 8a, also known as FEF, the frontal eye fields (Andersen et al. 1985; Andersen et al. 1990). The connection between LIP and FEF is reciprocal; the frontal eye fields send projections back to area LIP in a roughly topographic manner (Stanton et al. 1995). Area LIP also has reciprocal connections with multiple extrastriate visual areas, including areas V3, V3A, V4, and MT; the parieto-occipital visual area (PO); the dorsal prelunate area (DP); and area TEO, the occipital division of the intratemporal cortex (Blatt et al. 1990; Baizer et al. 1991; Webster et al. 1994). Parts of area LIP receive additional input from the anterior and medial divisions of intratemporal cortex (TEa and TEm), and the multimodal region of superior temporal cortex known as area IPa (Blatt et al. 1990). Within the parietal cortex, area LIP

sends a feedforward projection to area 7a, and interacts with areas MST, VIP and 7b (Blatt et al. 1990).

Subcortically, area LIP is interconnected with the lateral (non-retinotopic) pulvinar; this anatomical link further differentiates LIP from area 7a, which is connected to the medial pulvinar (Asanuma et al. 1985). Also in contrast to area 7a, area LIP sends a strong projection to the oculomotor (intermediate and deep) layers of the superior colliculus (Andersen et al. 1990; Baizer et al. 1993). Since these deeper layers of the superior colliculus project to the lateral pulvinar, the connection between area LIP and superior colliculus is indirectly reciprocal (Harting et al. 1980). Area LIP also sends projections to pontine nuclei linked to cerebellar oculomotor centers, and the pattern of this corticopontine projection strongly resembles that of the frontal eye fields (May and Andersen 1986).

Like many of the extrastriate visual areas to which it is anatomically connected, area LIP contains neurons with spatially tuned, oculocentric visual receptive fields (Gnadt and Andersen 1988; Barash et al. 1991b; Colby et al. 1995). Electrophysiological mapping studies have revealed a rough topography within area LIP; receptive field centers progress from the central to the peripheral visual field along the dorsal-ventral axis, and from the lower to the upper visual field along the anterior-posterior axis (Blatt et al. 1990). Unlike the receptive fields of neurons in area 7a, the receptive fields of neurons in area LIP tend to be relatively small and exclusively contralateral (Andersen et al. 1990; Blatt et al. 1990). Receptive field sizes are highly variable and increase with visual field eccentricity, but are typically 10–30° in diameter (Blatt et al. 1990; Barash et al. 1991b). In a recent study of LIP, Platt and Glimcher (1998) reported much smaller receptive fields; however, that study failed to control adequately for the presence of unresponsive cells, and therefore may have misidentified small receptive fields in noisy, untuned neural activity.

The close anatomical links between area LIP and oculomotor centers in the frontal eye fields and superior colliculus are evident in the responses of LIP neurons during saccadic eye movements. Many neurons in area LIP exhibit saccade-related activity in addition to (or instead of) visual activity (Mountcastle et al. 1975; Lynch et al.

1977; Hyvärinen 1982b). While saccade-related activity in area 7a tends to occur after eye movements have been initiated, activity usually begins before the saccade in area LIP, suggesting that area LIP is more involved in planning eye movements than area 7a (Andersen et al. 1990; Barash et al. 1991a).

The hypothesis that area LIP plays a role in planning eye movements to visual stimuli is supported by the results of microstimulation and lesion experiments. Electrical stimulation in area LIP and surrounding regions evokes saccadic eye movements (Shibutani et al. 1984; Kurylo 1991; Kurylo and Skavenski 1991; Thier and Andersen 1998; Mushiake et al. 1999), and lesions to area LIP and surrounding regions impair performance of voluntary delayed saccades (Quintana and Fuster 1993; Mazzoni 1994; Li 1996; Li et al. 1999). Similar saccadic deficits have been noted in humans, when trans-cranial magnetic stimulation is used to achieve a temporary disruption of processing in the posterior parietal cortex (Müri et al. 1996; Brandt et al. 1998).

Even more compelling evidence for involvement of area LIP in visual-to-oculomotor processing emerges from physiological studies of LIP activity in monkeys performing a visual memory-saccade task. In this task, the monkey holds his eye position on a fixation light while a visual stimulus is presented elsewhere in the visual field. The animal continues fixating after the stimulus disappears, and maintains fixation during a delay of several hundred milliseconds before the fixation light is extinguished. Then, the monkey makes a saccade to the remembered location of the visual stimulus presented earlier in the trial. Neurons in area LIP are active not only during the stimulus presentation and during the saccade, but also during the delay period of this task (Gnadt and Andersen 1988; Barash et al. 1991a; Barash et al. 1991b). This delay activity (also called memory activity) is thought to be related to the memory of the stimulus location or the plan to make a saccade.

Delay activity in area LIP, like visual and saccade-related activity, is strongly modulated by both eye position (Andersen et al. 1987; Andersen et al. 1990; Gnadt and Mays 1995) and head position (Brotschie et al. 1995; Snyder et al. 1998). These eye-position and head-position effects are thought to reflect the involvement of area LIP in transformation of retinocentric visual information into the head-centered and

body-centered frames of reference most appropriate for directing eye, head, and limb movements (Zipser and Andersen 1988; Andersen and Zipser 1988; Goodman and Andersen 1989; Salinas and Abbott 1996; Bremmer et al. 1998).

All these findings point to the conclusion that area LIP is involved in directing saccades to visual targets. However, the exact role of area LIP in visual-to-oculomotor processing has been the subject of long debate. In the first systematic physiological investigations of area 7, Mountcastle et al. (1975) interpreted activity in this region as primarily related to eye, head, and limb movements. Robinson et al. (1978), observing similar neural responses, argued that the region might be concerned with visual attention instead.

As explained above, more recent studies have generally concluded that this region forms an interface between sensory and motor systems; neurons in area LIP seem to have responses related both to spatial attention and to movement planning. Visual responses in area LIP are enhanced when an animal attends to a visual stimulus without looking at it (Colby et al. 1996); moreover, both the salience and the behavioral relevance of the visual stimulus influence activity in area LIP (Colby et al. 1996; Gottlieb et al. 1998). On the other hand, LIP neurons respond more strongly when the visual stimulus in the receptive field is a saccadic target than when the same stimulus is a visual distractor, even when the offset of the visual distractor is made relevant to the behavioral task (Platt and Glimcher 1997b). Buildup of activity in area LIP correlates with the probability that a stimulus will be a saccade target (Platt and Glimcher 1997a), and also with the monkey's certainty regarding the direction of an impending movement (Shadlen and Newsome 1996). Furthermore, when monkeys are required to make two sequential eye movements, activity in area LIP represents the next intended movement (Mazzoni et al. 1996a); this activity changes if the movement plan is suddenly altered (Bracewell et al. 1996). Finally, recent studies using both saccade and reaching tasks to analyze LIP activity have demonstrated that neurons in area LIP respond more strongly to visual stimuli which are targets for eye movements than to visual stimuli which are targets for arm movements (Snyder et al. 1997; Snyder et al. 1998). These findings demonstrate that visual responses in area

LIP contain a substantial motor component as well as a sensory component.

In another recent attempt to tease apart the sensory and motor components of neural activity in area LIP, Mazzoni et al. (1996b) trained monkeys to perform memory saccades to auditory as well as visual targets. Early studies of parietal cortex had found no responses to auditory stimuli in area 7 (Mountcastle et al. 1975; Hyvärinen 1982b; Koch and Fuster 1989). Therefore, it was thought that neurons in area LIP would not respond to auditory stimuli. Given this expected scenario, components of LIP activity most related to the intended movement might be distinguished from components of activity most related to the visual stimulus, through comparison of LIP responses recorded during auditory and visual memory-saccade trials.

The results of this experiment turned out to be even more interesting than was originally anticipated. During auditory memory-saccade trials, neurons in area LIP responded not only during the delay and saccade phases of the task, but also during the stimulus period (Mazzoni et al. 1996b). Responses to auditory stimuli were spatially tuned, strongly associated with visual responses, and comparable to visual responses in latency. Thus neurons in area LIP do respond to auditory stimulation, at least during an auditory-saccade task. These findings seem to contradict earlier studies, by suggesting that LIP neurons might exhibit sensory responses to auditory stimuli.

Auditory inputs to LIP clearly exist anatomically, but they are sparse and indirect compared to the visual inputs. The main auditory input to area LIP seems to come from temporo-parietal cortex (Tpt) in the superior temporal sulcus (Divac et al. 1977; Pandya and Kuypers 1969; Hyvärinen 1982b). Area Tpt is a high-level auditory association area thought to be involved in sound localization (Leinonen et al. 1980; Pandya and Yeterian 1985). As mentioned previously, the dorsal part of area LIP receives additional input from a multimodal area of superior temporal cortex called IPa (Blatt et al. 1990). Other regions of superior temporal cortex, including the superior temporal polysensory area (STP), also project to area LIP (Seltzer and Pandya 1991; Baizer et al. 1991) as well as to area 7a (Andersen et al. 1990).

Thus at least one auditory association area (area Tpt) and two multimodal regions

(IPa and STP) in the caudal superior temporal sulcus could be providing auditory input to area LIP. Auditory signals might also reach area LIP through feedback projections from frontal cortex or superior colliculus, both of which receive strong auditory projections (Kaas and Hackett 1998; Sparks and Hartwich-Young 1989). Whatever the source of input, however, it is clear from previous studies (Mazzoni et al. 1996b; Stricanne et al. 1996) that auditory signals reaching area LIP carry some information regarding the spatial location of auditory targets. Therefore, investigation of auditory responses in area LIP requires some basic understanding of sound localization.

1.3 Sound localization

The neurophysiology of sound localization is best understood in the barn owl (*Tyto alba*). The owl is an expert at sound localization within the animal kingdom, capable of localizing sounds with errors of less than 1° (Knudsen et al. 1979), compared to 2° for humans (Middlebrooks and Green 1991) and 4° for macaque monkeys (Brown et al. 1980; Brown et al. 1982). The neural pathways underlying sound localization have been mapped in greater detail for the owl than for any other animal. Therefore, sound localization in primates will be explained here in comparison with sound localization in the owl. Except where noted otherwise, these summaries are based on comprehensive reviews by Knudsen et al. (1987), Konishi et al. (1988), Konishi (1991), Middlebrooks and Green (1991), Konishi (1992), Brugge (1992), Webster (1992), and Konishi (1993).

1.3.1 Barn owls

Owls compute sound location from two cues: differences in the phase delay of sounds arriving at the two ears (interaural time differences, or ITDs), and differences in sound intensities at the two ears (interaural intensity differences, or IIDs). Interaural time differences provide azimuth cues, because sounds coming from one side of the animal arrive at the ear on that side first. In contrast, interaural intensity differences

provide elevation cues. The external ears of owls are asymmetric; the right ear is angled upward, and the left ear downward. Therefore, the right ear is more sensitive to sounds coming from above the animal's head, and the left ear more sensitive to sounds coming from below the head.

Separate neural pathways in the owl's midbrain extract ITD and IID cues from auditory nerve discharges. Primary auditory fibers fire action potentials which are phase-locked to particular sound frequencies (and time-locked to the moment of sound arrival at the ear); moreover, the overall discharge rate is a monotonic function of sound intensity. Thus the auditory nerve carries multiplexed timing and intensity signals. These signals are separated at the level of the cochlear nuclei, where the auditory nerve bifurcates. One branch enters the nucleus magnocellularis, which is specialized for processing timing information. The other branch enters the nucleus angularis, which extracts overall firing rate information. These two cochlear nuclei form the first stages of parallel pathways for computing ITD and IID.

The first station in the ITD pathway after nucleus magnocellularis is nucleus laminaris, which cross-correlates inputs from the two ears to extract the times of binaurally coincident auditory fiber discharges within each frequency band. These times represent frequency-specific interaural phase differences, not interaural time differences, because phase-locked impulse trains from the two ears may be offset by multiple phase cycles. This ambiguity in the timing information is removed through comparison of interaural phase differences across frequency bands. The extraction of ITD from interaural phase differences across multiple frequency channels occurs in the central nucleus of the inferior colliculus (ICc), which receives direct input from nucleus laminaris (and also indirect input, through the nucleus ventralis lemnisci lateralis, pars anterior).

The IID pathway starting from nucleus angularis courses through the nucleus ventralis lemnisci lateralis pars posterior, in which signals from the two ears are compared. Inhibitory inputs from one ear are combined with excitatory inputs from the other, to extract interaural intensity difference information within each frequency band. These frequency-specific IID signals are combined across frequency channels

in the central nucleus of the inferior colliculus, and further processed to obtain final IID estimates.

Thus the ITD and IID pathways converge in the ICc, where neurons selective to combinations of interaural time and intensity differences first emerge. The convergence culminates with formation of a map of auditory space in the external nucleus of the inferior colliculus (ICx). This auditory space map in the inferior colliculus projects onto the optic tectum, which also contains a map of visual space; the two sensory maps become aligned during development (Knudsen and Brainard 1995). Within the optic tectum, the joint auditory-visual sensory space map projects to a motor map controlling orienting movements.

Recently, another sound localization pathway has been identified in the owl, involving forebrain rather than midbrain structures (Knudsen et al. 1995; Cohen and Knudsen 1995; Cohen et al. 1998). This forebrain pathway diverges from the midbrain pathway at the central nucleus of the inferior colliculus. Auditory information in ICc is sent not only to ICx along the midbrain pathway, but also to the auditory thalamus (nucleus ovoidalis), the gateway to the forebrain pathway. The auditory thalamus projects to Field L, a primary auditory area in the owl forebrain. Neurons in Field L have binaural tuning properties similar to those of neurons in ICc, but are arranged into a fractured topography instead of a space map (Cohen and Knudsen 1998). Continuing along the forebrain pathway, Field L projects to the auditory archistriatum (AAr), a newly discovered forebrain region. Neurons in AAr are sharply tuned for sound source location (Cohen and Knudsen 1995) and also involved in gaze control (Knudsen et al. 1995). Thus, like the optic tectum, AAr contains representations of both auditory space and orienting movements. However, unlike the optic tectum, AAr representations of auditory space and orienting movements form locally organized clusters but not continuous maps.

The sharp tuning of AAr neurons for binaural spatial cues depends only on input from Field L, not on input from the auditory space map in ICx (Cohen et al. 1998). Therefore, the representation of auditory space in the forebrain is independent of the midbrain pathway. The independence of the two pathways has been confirmed

in lesion studies. After inactivation of either the forebrain pathway or the midbrain pathway alone, owls can still orient to auditory stimuli; the ability to localize sounds is lost only when both pathways are disrupted (Knudsen et al. 1993). Thus the forebrain and midbrain pathways form parallel sound localization systems. Cohen et al. (1998) have proposed that “the forebrain pathway primarily participates in voluntary shifts of gaze, such as those that require access to memory stores, and... the midbrain pathway... is particularly important for short latency, reflexive orienting movement.” This hypothesis is supported by recent behavioral studies, which demonstrate that after inactivation of AAr, owls can still orient and fly toward an ongoing sound, but they are no longer able to orient to the remembered location of a previously presented auditory target (Knudsen and Knudsen 1996).

1.3.2 Primates

Neural pathways for sound localization in primates appear to be similar to those in owls, but primates exploit sound localization cues in a slightly different way. Since primates do not have vertically asymmetric ear openings, IIDs provide information about azimuth rather than elevation. However, IIDs are useful azimuth cues only for high-frequency sounds, because the head induces a significant sound shadow only if the sound wavelength is small relative to the diameter of the head. Fortunately, ITDs are most useful as azimuth cues at low frequencies, because interaural phase differences correspond directly to interaural time differences for long-wavelength sounds (i.e., there is no phase ambiguity to be resolved). Primates therefore use both ITDs and IIDs as azimuth cues, but over different frequency ranges; for humans, the relevant ITD-IID switchover point is about 3 kHz.

Primates extract sound elevation information from another spatial cue: the spectral envelope of the sound after it has been filtered by the external ear. Primate ear flaps, or *pinnae*, are highly convoluted and vertically offset with respect to the ear canal. These properties ensure that the pinnae impose upon sounds arriving at the ear a spectral filter which varies with the elevation of the sound source. Sounds

reaching the ear canal directly are combined with delayed versions of the same signal reflected off the convolutions of the pinnae. The sound which ultimately arrives at the eardrum has a comb-filtered spectrum containing peaks and notches characteristic of a particular source elevation. In contrast to ITD and IID cues, these spectral cues can be monaural, in that they do not depend on comparisons between the ears. Therefore, primates can localize sounds in elevation with only one ear.

Analysis of spectral cues does have drawbacks as a sound localization strategy, however. Useful spectral cues are available only when the sound incident upon the ear contains energy at high frequencies, and is relatively broad in bandwidth. Primates are therefore very inaccurate at judging elevation of low-frequency noises and pure tones. More significantly, analysis of spectral cues requires prior assumptions regarding both the likely spectrum of the sound source and the exact filtering properties of the pinnae. Vertical localization in primates is therefore very sensitive to manipulations of the source spectrum and modifications of the pinnae. Interestingly, primates are adept at recalibrating elevation judgments; even adult humans can quickly relearn mappings between spectral cues and sound elevation when the convolutions of the pinnae are occluded (Hofman et al. 1998).

The neural pathways for processing auditory spatial information in primates appear to be very similar to those of owls (although the neural substrates for the spectral analysis inherent in primate elevation judgments are not yet known). The medial superior olive has been identified as the mammalian homologue of the avian nucleus laminaris; this structure computes frequency-specific interaural phase differences in the midbrain ITD pathway. Along the midbrain IID pathway, the mammalian homologue of the avian lateral lemniscal nuclear complex is the lateral superior olive, which extracts frequency-specific interaural intensity differences. As in the owl, the ITD and IID streams converge in the inferior colliculus. However, in primates and other mammals, a complete map of auditory space does not emerge in the inferior colliculus, but only in the primate homologue of the avian optic tectum: the superior colliculus. Like the optic tectum, the superior colliculus contains sensory maps of visual and auditory space, along with motor maps of orienting movements.

Forebrain pathways for processing auditory space have not been characterized fully for primates, or indeed for any mammal. However, there is very likely to be a homologue of the owl forebrain pathway in primates. Some parallels are already evident; for example, the functional organization of avian Field L resembles that of primate primary auditory cortex (Cohen and Knudsen 1998). Moreover, the gaze fields of the owl's forebrain archistriatum (including AAr) are both anatomically and physiologically similar to the primate frontal eye fields (Knudsen et al. 1995). Representations of auditory space in FEF are poorly understood, and so it is not yet known if a direct homologue of AAr exists within FEF. However, neurons in FEF do have spatially tuned responses to auditory stimuli (Vaadia et al. 1986; Schall 1991a; Russo and Bruce 1994), and FEF receives strong projections from multiple auditory association areas (Pandya and Yeterian 1985; Romanski et al. 1999). There are probably several stages of auditory spatial information processing between primary auditory cortex and FEF, but understanding of auditory association areas in the primate is still rudimentary. Auditory association areas thought to be involved in sound localization include area Tpt (Leinonen et al. 1980; Kaas and Hackett 1998) and caudomedial area CM (Rauschecker et al. 1997).

Like the forebrain sound localization pathway in the owl, the putative forebrain pathway in the primate may differ from the midbrain sound localization pathway in being specialized for voluntary rather than reflexive orienting movements to auditory targets. In fact, this type of functional specialization of midbrain and forebrain localization circuits has already been documented in primates for visually guided saccadic eye movements, as explained further below.

1.4 Saccadic eye movements

Fast eye movements to visual targets seem to be controlled by two interconnected but relatively independent neural pathways: a midbrain pathway dominated by the superior colliculus, and a forebrain pathway involving the frontal eye fields. These pathways are briefly outlined below; for details beyond the scope of this quick sum-

mary, see reviews by Sparks and Hartwich-Young (1989), Goldberg and Segraves (1989), and Robinson and McClurkin (1989).

The superior colliculus receives extensive visual inputs, both directly from the retina and striate cortex (superficial layers), and indirectly from association areas in the occipital, temporal, parietal, and frontal lobes (deep layers). The deep layers of the colliculus send ascending projections back to cortex, and also descending projections to brainstem saccade generators. Although there is little evidence for direct connections between the superficial and deep layers, the deep layers may receive indirect input from the superficial layers via the pulvinar, thalamus, and cortex. Electrical stimulation in the deep layers of the superior colliculus evokes saccades and other orienting movements.

The frontal eye fields also receive extensive visual inputs, primarily from high-level extrastriate areas in parietal cortex (including area LIP). Neurons in FEF project heavily to the superior colliculus, especially to the deeper layers; thus eye movements initiated in FEF are mediated in part by activation of superior colliculus. However, FEF also projects directly to brainstem saccade generators. Microstimulation in the frontal eye fields evokes saccades, even when the superior colliculus has been inactivated (Schiller 1977).

Monkeys lose the ability to make saccades if both the superior colliculus and the frontal eye fields are lesioned, but they are still able to saccade if only one of the two structures is ablated (Schiller et al. 1980). Saccadic deficits after superior colliculus lesions include minor but persistent reductions in saccade frequency, accuracy, latency, and velocity (Schiller et al. 1980; Schiller et al. 1987). Deficits resulting from inactivation of the frontal eye fields are much more subtle. Visually guided saccades are virtually unaffected by FEF lesions; however, deficits in predictive saccades, memory saccades, and saccade target selection have been reported (Bruce and Borden 1986; Deng et al. 1986; Schiller and Chou 1998).

These findings seem to indicate that the collicular gaze control pathway in primates is most critical for reflexive eye movements to visual targets, while the cortical pathway is specialized for control of voluntary, predictive, and learned eye move-

ments. Assuming that these results hold for auditory saccades, the parallels between the primate and owl systems may be extended to include the roles of midbrain and forebrain pathways in auditory orienting.

However, there is one critical difference between owls and primates, which limits the usefulness of the cross-species comparison in the context of saccadic eye movements. Owls have an extremely limited oculomotor range, and therefore orient to sounds primarily through head movements. Since auditory localization cues are head-centered (because the ears are fixed in the head), the owl's sensory and motor reference frames are always aligned during localization of auditory targets. In contrast to owls, primates have an extremely large oculomotor range; the eyes can deviate up to 40 degrees while the head remains fixed. When the eyes rotate in the head, head-centered and eye-centered reference frames become misaligned. Eye movements to auditory stimuli in primates must therefore involve transformation of head-centered auditory signals into eye-centered movement commands. The side effects of this transformation are evident in many aspects of primate behavior and neurophysiology.

1.5 Sound localization and eye movements

Psychophysical, behavioral, and neurophysiological evidence suggests that sound localization is critically dependent on eye movements in primates (for reviews, see Welch and Warren 1986; King 1988). Indeed, sound localization itself, even in the absence of saccades to auditory targets, is influenced by eye position. Sustained deviation of the eyes induces a shift in the perceived auditory midline toward the direction of gaze (Thurlow and Kerr 1970; Lackner 1973; Lewald and Ehrenstein 1996b; Lewald and Ehrenstein 1996a; Lewald 1998). Moreover, sound localization is more accurate when subjects view an illuminated, textured background than when localization tasks are performed in the dark, apparently because internal estimates of eye position are more accurate in the light than in the dark (Platt and Warren 1972; Mastroianni 1982). Accuracy of sound localization in the dark improves when subjects are allowed to make saccades to the perceived location of the auditory target, suggesting

that eye movements are used to stabilize auditory position memory (Jones and Kabanoff 1975). Finally, localization of auditory targets with saccadic eye movements seems to be strongly influenced by initial eye position. Both the accuracy and the latency of saccades to auditory targets are dependent not on the absolute azimuth of the auditory target with respect to the head, but on the eccentricity of the auditory target with respect to the eyes (Jay and Sparks 1990; Zambarbieri et al. 1995; Yao and Peck 1997). These findings suggest that sound localization and eye movement systems are intertwined in the primate brain.

Neurophysiological interactions between the sound localization and eye movement systems have been observed in the superior colliculus and the frontal eye fields. In both areas, responses to auditory stimuli during saccades to auditory targets are strongly influenced by initial eye position (Jay and Sparks 1987a; Russo and Bruce 1994). In other words, auditory signals in both areas seem to have been transformed (at least partially) from the head-centered reference frame of the two ears, into the eye-centered reference frame appropriate for directing saccades.

Recent studies have demonstrated that neurons in area LIP also tend to respond to auditory signals in an eye-centered reference frame. Stricanne et al. (1996) recorded LIP responses to auditory stimuli while monkeys performed an auditory memory-saccade task from different initial fixation positions. While a few of the cells recorded in area LIP appeared to have head-centered auditory receptive fields, the vast majority of responses to auditory stimuli in area LIP were significantly modulated by eye position. These findings suggest that area LIP, the frontal eye fields, and the deep layers of the superior colliculus all receive auditory signals which have been partially or completely transformed into an eye-centered reference frame.

The results also deepen the mystery regarding the exact role of LIP in auditory-to-oculomotor transformations. Given its intermediate placement between sensory and motor regions of the brain, area LIP might be receiving head-referenced auditory signals from auditory association areas, transforming those signals into the eye-centered reference frame most appropriate for planning eye movements, and then sending this information on to the superior colliculus and frontal cortex. Alternatively, area LIP

might be receiving auditory signals via feedback from superior colliculus and frontal eye fields; therefore, it is possible that the auditory information is transformed into an eye-centered reference frame well before it reaches LIP.

1.6 Responses to auditory stimuli in area LIP

Given that area LIP is involved in oculomotor planning, that LIP neurons have spatially tuned responses to auditory stimuli, and that sound localization in primates is strongly influenced by oculomotor behavior, investigations of auditory responses in area LIP may clarify not only the role of area LIP in auditory-to-oculomotor processing, but also LIP function and auditory-to-oculomotor transformations in general. This dissertation presents experiments which were designed to investigate the effects of training and behavioral context on LIP responses to auditory stimuli, while comparing auditory responses to visual responses. Chapters 2, 3, and 4 directly address the following questions:

- Are responses to auditory stimuli in area LIP dependent on auditory-saccade training?
- Are auditory responses in area LIP affected by the behavioral context in which auditory stimuli appear?
- Are responses to auditory stimuli in area LIP qualitatively different from responses to visual stimuli?

Ultimately, the results of the experiments not only answer these questions, but also provide important insight into three much broader issues:

- What is the role of area LIP in auditory-to-oculomotor processing?
- What is the function of area LIP overall?
- What is the neurophysiological basis for auditory-to-oculomotor transformation in primates?

The implications of the results with respect to these larger issues are discussed in the final chapter.

Chapter 2 Training Effects

The lateral intraparietal area, a region of macaque posterior parietal cortex, has long been considered unresponsive to auditory stimulation. However, recent reports indicate that neurons in this area respond to auditory stimuli in the context of an auditory-saccade task. To what extent are responses to auditory stimuli in area LIP dependent upon auditory-saccade training? To address this question, recordings were made from area LIP in two monkeys, both before and after the animals had been trained to make saccades to auditory stimuli. At the outset of the study, the animals were naive with respect to the auditory tasks, but had been previously trained to perform visual tasks. Both before and after auditory-saccade training, neural activity was recorded in area LIP while the animals were performing a fixation task involving presentations of auditory and visual stimuli. Among 172 LIP neurons recorded before auditory-saccade training, the number of cells responding to auditory stimuli did not reach significance, but about half the neurons responded to visual stimuli. An information theoretic analysis, used to quantify the degree of spatial tuning in auditory and visual responses, confirmed that the firing rates of neurons recorded in area LIP before auditory-saccade training carried significant amounts of information about visual stimulus location, but no information about auditory stimulus location. After auditory-saccade training, however, 12% of 160 cells recorded in LIP responded to auditory stimuli, while the proportion of neurons responding to visual stimuli remained about the same as before training. Firing rates of these 160 neurons recorded after auditory-saccade training carried significant amounts of information about auditory as well as visual stimulus location. Auditory-saccade training therefore generated responsiveness to auditory stimuli *de novo* in area LIP. The results indicate that some LIP neurons can become responsive to auditory stimuli even in a passive fixation task, once the animals have learned that these stimuli are important for oculomotor behavior.

2.1 Introduction

As explained in Chapter 1, the lateral intraparietal area of macaque posterior parietal cortex participates in the intermediate stages of visual-to-oculomotor transformations. Located in the middle of the dorsal visual stream, the “where” pathway in vision (Ungerleider and Mishkin 1982), area LIP receives strong visual inputs from multiple extrastriate visual areas, and is interconnected with oculomotor centers in the frontal cortex (Andersen et al. 1985; Andersen et al. 1990; Blatt et al. 1990; Stanton et al. 1995), the superior colliculus (Lynch et al. 1985), and the cerebellum (via the pontine nuclei; May and Andersen 1986). This anatomical evidence indicates that area LIP is involved in conversion of visual input to oculomotor output (Andersen 1987; Gnadt and Andersen 1988; Colby et al. 1996).

Like the anatomy, the physiology of LIP suggests that this area links visual processing with oculomotor planning. Neurons in area LIP are activated during visual stimulation (Blatt et al. 1990), during visual attention (Colby et al. 1996; Gottlieb et al. 1998), during eye movement planning (Gnadt and Andersen 1988; Mazzoni et al. 1996a; Bracewell et al. 1996; Shadlen and Newsome 1996; Platt and Glimcher 1997b), and during eye movements (Mountcastle et al. 1975; Lynch et al. 1977; Hyvärinen 1982b; Barash et al. 1991a). Visual responses in area LIP are spatially tuned in an oculocentric coordinate frame (Gnadt and Andersen 1988; Barash et al. 1991b; Colby et al. 1995) and additionally are modulated by eye position (Andersen et al. 1990). Neurons in area LIP respond more strongly when the visual stimulus in the receptive field is a saccadic target than when the same stimulus is a visual distractor, even when the offset of the visual distractor is made relevant to the behavioral task (Platt and Glimcher 1997b). Moreover, activity in area LIP seems to follow the eye movement plan (Mazzoni et al. 1996a; Bracewell et al. 1996), and LIP neurons respond more strongly to visual stimuli which are targets for eye movements than to visual stimuli which are targets for arm movements (Snyder et al. 1997; Snyder et al. 1998). These findings indicate that area LIP plays a special role in directing eye movements to visual stimuli.

Since auditory as well as visual stimuli can serve as targets for eye movements, area LIP could conceivably be involved in auditory-to-oculomotor as well as visual-to-oculomotor transformations. Although the known auditory inputs to LIP are sparse compared to the visual inputs, at least one auditory association area (area Tpt) is clearly linked to the posterior parietal region (Pandya and Kuypers 1969; Divac et al. 1977; Hyvärinen 1982b). Polysensory areas in the superior temporal sulcus also project directly to the intraparietal sulcus (Blatt et al. 1990; Seltzer and Pandya 1991; Baizer et al. 1991). Moreover, movement-related auditory responses have been observed in several regions of the brain which are anatomically connected to area LIP, including the frontal eye fields (Vaadia et al. 1986; Russo and Bruce 1994) and the deep layers of the superior colliculus (Jay and Sparks 1987a).

Early physiological investigations of LIP and surrounding regions found no auditory activity in this area (Mountcastle et al. 1975; Hyvärinen 1982b; Koch and Fuster 1989). More recently, however, Mazzoni et al. (1996b) and Stricanne et al. (1996) recorded responses to auditory stimulation in area LIP in the context of an auditory memory-saccade task. Monkeys were trained to remember the location of an auditory stimulus, and to make a saccade to the remembered location after a delay. Neurons in area LIP were active not only during the movement and delay phases of this task, but also during the auditory stimulus presentation (Mazzoni et al. 1996b; Stricanne et al. 1996). The presence of activity during the stimulus presentation period suggests that auditory responses in LIP might be sensory in nature. Thus these recent results, which show that neurons in area LIP respond to auditory stimuli during an auditory-saccade task, seem to contradict the earlier studies, which reported no evidence for activity in area LIP during auditory stimulation.

There are four possible explanations for this apparent discrepancy. First, early studies may simply have failed to detect auditory responses in LIP, because the search for such responses was not conducted in a systematic fashion. Second, LIP neurons may respond to auditory stimuli only after auditory-saccade training. Third, auditory responses may appear in LIP only when the animal is engaged in an auditory-saccade task. And fourth, responses to auditory stimuli in area LIP may be dependent both

on auditory-saccade training and on the immediate behavioral task the animal is performing. This last possibility stands in contrast to the second and third possibilities, which propose either that auditory responses in area LIP arise through long-term training and are unaffected by immediate behavioral context, or that auditory responses are entirely dependent on task and unaffected by training history.

This chapter investigates the effects of training on auditory responses in area LIP, by analyzing data recorded while animals were performing the same behavioral task before and after auditory-saccade training. The analyses demonstrate both that area LIP is unresponsive to auditory stimulation before auditory-saccade training, and that responses to auditory stimuli emerge in LIP after training. Chapter 3 examines how auditory and visual responses observed after auditory-saccade training are affected by immediate behavioral context, by comparing responses recorded in a fixation-only task with responses recorded during an auditory-saccade task. Taken together, these two chapters show that both training history and behavioral task have a strong influence on the auditory responsiveness of area LIP. Reports of these results have been published in abstracts (Linden et al. 1996; Linden et al. 1997) and in an article (Grunewald et al. 1999).

2.2 Methods

2.2.1 Animals and animal care

Two male rhesus monkeys (*Macaca mulatta*) were used in this study. Neither had participated in any previous auditory experiments. Monkey B was 6 years old at the beginning of this study, and had previously participated in experiments involving visually-guided eye movements. At the conclusion of the present study, monkey B was euthanized; all histological data shown in this chapter are from this monkey. Monkey Y, who was not involved in any experiments prior to this study, is still a subject in ongoing research.

2.2.2 Surgical procedures

All surgical and animal care procedures were in accordance with National Institutes of Health guidelines, and were approved by the California Institute of Technology Institutional Animal Care and Use Committee. Surgeries were performed under sterile conditions, using general anesthesia (10 mg/kg sodium pentobarbital intravenously). Heart rate, respiration rate, and temperature were monitored throughout each surgery. At the start of the experiments, a stainless-steel head post and a dental acrylic head cap were implanted onto the skull of each animal. In the same procedure, a scleral search coil was implanted to monitor eye movements (Mountcastle et al. 1975; Judge et al. 1980). Analgesics and systemic antibiotics were administered for several days after surgery, and the animals were allowed to recover for at least a week before any behavioral training began.

After recovery, the animals were trained to perform a fixation task. Since fixation of a visual target involves a saccade to acquire the target, the animals were implicitly trained to perform visually-guided eye movements. Once the animals performed the fixation task with sufficient accuracy (approximately 95% success rate), a second surgery was performed in which a stainless steel recording chamber was mounted over posterior parietal cortex normal to the skull surface (stereotaxic coordinates at center: 6 mm posterior, 12 mm lateral). The chamber was implanted over the left hemisphere in monkey B, and over the right hemisphere in monkey Y.

2.2.3 Experimental setup

All experiments and behavioral training sessions were conducted in complete darkness, in a double-walled sound-attenuating anechoic chamber (Industrial Acoustics Company, Inc.). The walls of the chamber attenuated external sounds above 200 Hz by at least 60 dB, and the interior of the chamber qualified as anechoic for sounds between 200 Hz and 16 kHz (inverse-square-law test deviations from theoretical free-field conditions of less than 1.0 dB in 500 Hz – 8 kHz frequency range and less than 1.5 dB in 200 Hz – 16 kHz frequency range; see Schmitt 1983 for further explanation).

While inside the chamber, the monkey was monitored continuously with an infrared camera and a microphone. The monkey's primate chair was mounted inside a frame of 90-cm diameter magnetic coils used to measure eye movements.

The monkey's head was held fixed during all behavioral training and recording sessions; all stimulus locations are therefore specified relative to the center of the monkey's head, in degrees azimuth right or left of the median sagittal plane and in degrees elevation above or below the visual plane. Two fixed arrays of speakers and light-emitting diodes (LEDs) were used to present auditory and visual stimuli. An LED was mounted at the center of each speaker in the array. In the earliest experiments, a hexagonal array was used; this array was replaced with an improved rectangular version in later experiments. The hexagonal stimulus array consisted of 19 speaker/LED devices arranged hexagonally, such that the center-to-center separation of the stimuli was 12 degrees. The rectangular array was concave and consisted of 25 speaker/LED devices, with a minimum center-to-center spacing of 8 degrees. The stimulus array was positioned 80 cm from the monkey's head, and the central speaker/LED coincided roughly with the straight-ahead eye position of the monkey. Both the stimulus array and the magnetic coil frame were padded with sound-absorbing acoustical foam (Sonex) to dampen echoes from their surfaces.

Visual stimuli were 500 ms flashes of 70 cd/m² red light from the LEDs, which each subtended 0.4 degrees. Free-field auditory stimuli were 500 ms bursts of band-limited noise (5–10 kHz, 5 ms rise/fall times, 70 dB SPL). This noise band was chosen because macaque monkeys have been reported to localize 5–10 kHz bandlimited noise well in azimuth (Brown et al. 1980), and because the frequency responses of the speakers were relatively flat (± 10 dB SPL) within this range. In the recording experiments before the animals had been trained to perform auditory saccades, the speakers sounded about the same, but their responses had not been equalized. For saccade training and post-training recordings, the input to each speaker was adjusted to equalize the output amplitude spectrum to 70 ± 2 dB SPL within the 5–10 kHz frequency band, as measured at the location where the monkey's head would be during an experiment. This equalization was performed to ensure that the monkeys were

performing a localization task, rather than a spectral recognition task, when they were instructed to perform eye movements. Chapter 3 addresses the localization task in greater detail.

2.2.4 Behavioral paradigms

Each monkey participated in three phases of the experiment: the *pre-training* phase (before auditory-saccade training), the *training* phase (during auditory-saccade training), and the *post-training* phase (after auditory-saccade training). Note that in the pre-training phase the two monkeys were not completely naive, since they were both familiar with the fixation task and with visual saccade tasks used for eye position calibration. The behavioral phase name therefore designates the animal's training with respect to auditory saccades alone.

Pre-training

In the pre-training behavioral task, the monkey fixated the central LED in the hexagonal or rectangular array while auditory or visual stimuli were briefly presented at other locations in the array. Each trial began with presentation of the fixation light. The monkey acquired fixation, held fixation for 1000–1500 ms until an auditory or visual stimulus appeared, continued fixating during the 500 ms stimulus presentation period, and then maintained fixation for another 500 ms after stimulus offset. This task is illustrated in Figure 2.1. The monkey was required to hold his eye position within a square of half-width 2–3° (earlier experiments) or a circle of radius 2–3° (later experiments) centered on the fixation light throughout the trial. If the animal succeeded in this task, he was rewarded with a small quantity of water or juice. Auditory and visual trials involving stimuli at different locations were all randomly interleaved.

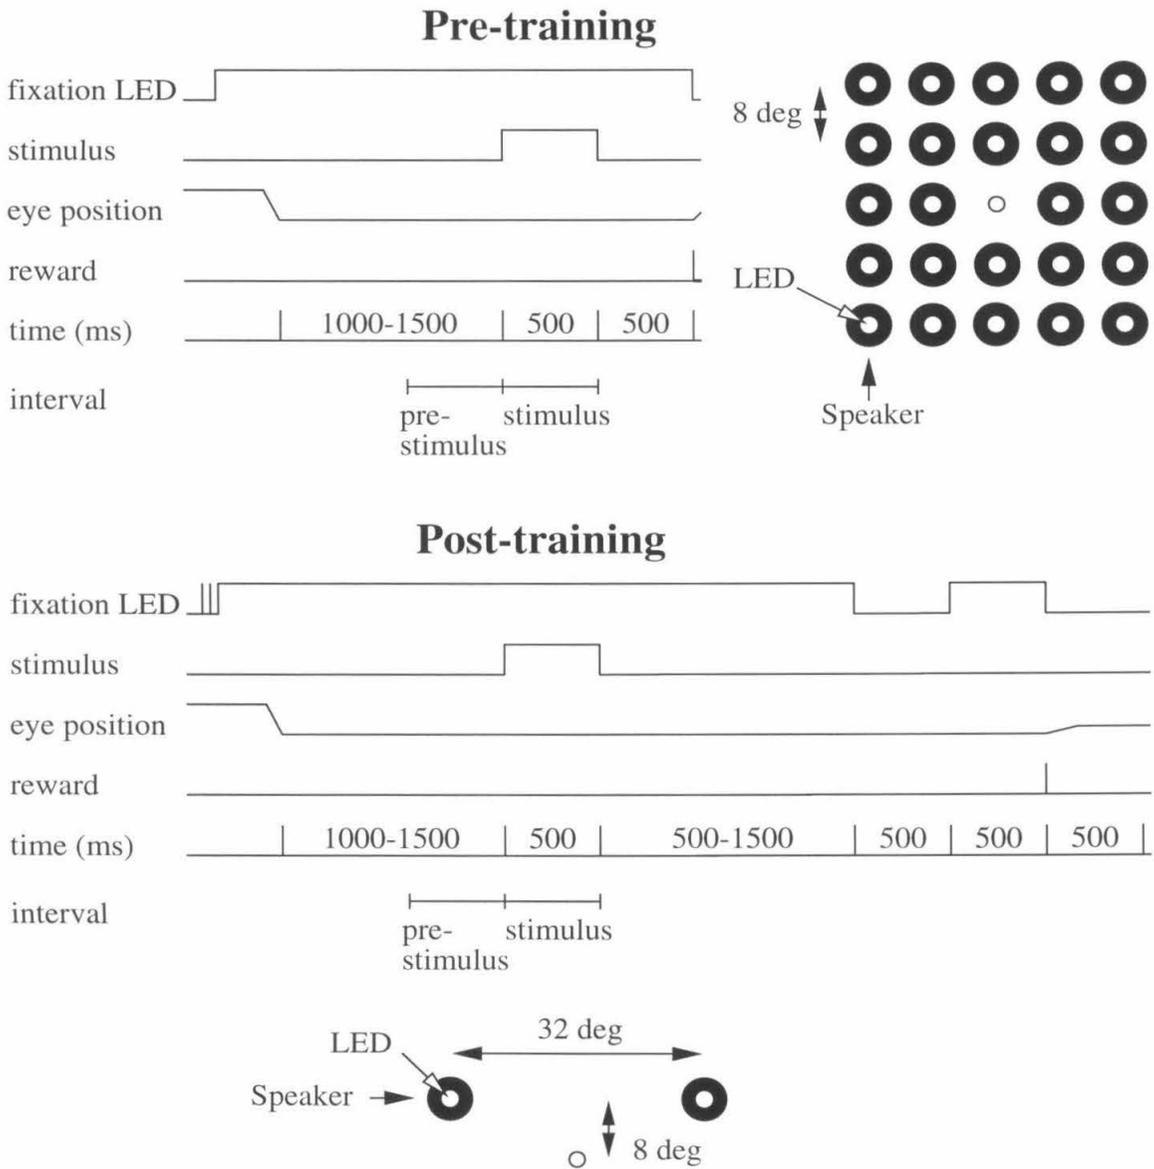


Figure 2.1: Diagrams of the pre- and post-training fixation tasks and stimulus arrays. In both the pre- and post-training tasks, the monkey held fixation for 1000–1500 ms before stimulus onset, maintained fixation through the stimulus interval, and continued fixating for at least 500 ms after stimulus offset. The behavioral requirements of the two tasks were therefore identical in the pre-stimulus and stimulus intervals. In the post-training fixation task, the monkey was also required to maintain fixation in darkness after fixation light offset, to ensure that no late saccades were performed. See text for details.

Training (auditory saccades)

Once pre-training data collection was completed, the animals were trained to perform saccades to auditory targets. Surprisingly, the auditory-saccade task was not easy for the monkeys to learn. Initial attempts to train auditory saccades without the use of visual feedback were not successful; training was ultimately accomplished by presenting an auditory stimulus, requiring the animal to complete a saccade to the auditory stimulus in darkness, and then presenting a visual stimulus at the target location and requiring a corrective saccade. The visual feedback stimulus never appeared simultaneously with the auditory stimulus. Eye movements to auditory targets were deemed to have sufficient accuracy during training if they ended within a circular window of radius 16° around the stimulus. A reward was administered if the auditory saccade was accurate, and if the subsequent corrective saccade was also accurate. The visual feedback stimulus was gradually moved further back in time, so that eventually it appeared 500 ms after the animal had successfully acquired the auditory target window. In all trials, however, a visual feedback stimulus appeared at the end of the trial. This practice was continued even during the recording sessions.

During training and in the subsequent experiments, only the rectangular stimulus array was used. After seven months of training, monkey B had learned to perform eye movements to four targets spaced in azimuth. However, the accuracy of the eye movements was not very high (about $\pm 12^\circ$), so for the recording experiments only two target locations were used: $(-16^\circ, +8^\circ)$ and $(+16^\circ, +8^\circ)$. Monkey Y did somewhat better; after 5 months of training, he had learned to make auditory saccades to 9 targets with the same accuracy. However, to maintain consistency across the two animals, the same two target locations were employed for both animals during the recording experiments.

When the monkeys had learned to perform auditory saccades to within 12° of the target locations with a success rate exceeding 80%, auditory and visual memory-saccade training began. In the memory-saccade task, the monkeys were required to maintain fixation through presentation of an auditory or visual stimulus at one

of the two target locations, and then to continue fixating for a delay period after stimulus offset. Once the fixation light had been extinguished, the animals had to make a saccade to the remembered location of the auditory or visual stimulus. Both monkeys learned the memory-saccade paradigm for auditory stimuli in a single day, but training of visual memory-saccades required a few days.

Post-training

Once the animals had learned the memory-saccade task, they were trained to perform a modified version of the pre-training fixation task. In this modified fixation task, the fixation light flashed twice before staying on. The flash sequence indicated to the monkey that a fixation trial was about to occur rather than a memory-saccade trial (see Chapter 3). The animal was required to hold his eye position within $2\text{--}3^\circ$ of the fixation light. After a variable interval (1000–1500 ms), an auditory or visual stimulus appeared for 500 ms at one of the two locations used in the training phase and the memory-saccade task: $(-16^\circ, +8^\circ)$ or $(+16^\circ, +8^\circ)$. The fixation light remained on, and the monkey continued fixating, through the stimulus presentation and for 500–1500 ms following the disappearance of the stimulus. Then, the fixation light was extinguished, but the monkey continued fixating at the same location in darkness for 500 ms. Continued fixation was required to ensure that the monkey did not make a saccade to the stimulus location immediately after fixation light offset. Finally, the fixation light reappeared, and if the monkey fixated it for an additional 500 ms he was rewarded with a drop of water or juice. Eye position was monitored for at least 500 ms following the reward, again to ensure that the monkey did not make a saccade to the stimulus location. Blocks of fixation trials were alternated with blocks of memory-saccade trials (described in Chapter 3); within each task block, trials involving auditory and visual (and left and right) stimuli were randomly interleaved.

2.2.5 Recording procedures

Single-unit extracellular recording was performed using tungsten microelectrodes and a hydraulic microdrive (Frederick Haer & Co.). All penetrations were approximately normal to the gyral surface. To help ensure that recordings came from area LIP (within the intraparietal sulcus) rather than area 7a (on the gyral surface), the electrode was advanced to 2500–3000 μm below the dura at the start of each recording session. A guide tube protected the electrode during penetration of the dura and served as the reference input for the differential microelectrode amplifier. Electrode impedances were typically 0.5–2.0 $\text{M}\Omega$ at 1 kHz. The electrode signal was amplified by a factor of 200,000, band-pass filtered (Krohn-Hite) between 600 Hz and 5 kHz, and monitored continuously on an oscilloscope and audio monitor. Single units were isolated using a variable-delay voltage-time window discriminator (Tucker-Davis Technologies), and times of spike occurrence were recorded with 1 ms accuracy.

Eye position was monitored with a precision of $\pm 0.1^\circ$ using the scleral search coil technique (Mountcastle et al. 1975; Judge et al. 1980), and was recorded at 1000 samples/sec. Each behavioral training or recording session began with a calibration of eye position recording equipment, during which the animal fixated visual stimuli at various locations on the stimulus array.

2.2.6 Recording strategies

Data were collected in only two of the three behavioral phases: the pre-training phase, and the post-training phase. Recording strategies used in each phase are described below.

Pre-training

In pre-training experiments, the electrode was advanced while the monkey was performing the pre-training fixation task described previously. Any neuron which was encountered, and which could be kept isolated long enough to characterize, was included in the pre-training database. In other words, there was no bias in the recording

strategy that might have favored finding cells with auditory or visual responses.

Post-training

For post-training data collection, a slightly different recording strategy was used. While the electrode was advanced in search of neurons, the monkey was performing interleaved auditory and visual memory-saccade trials. Data were recorded from any isolated cell which appeared to exhibit a response during any period of either the auditory or the visual memory-saccade task. Thus there were two differences between the pre- and post-training recording strategies: the task the animals were performing during the search for neurons, and the cell selection criteria. The post-training recording strategy resulted in a bias favoring neurons which responded during the memory-saccade task. Analyses and controls presented in the Results section indicate that the findings of this chapter are not affected by these differences. This issue is also addressed in the Discussion section.

2.2.7 Database

Pre-training

The pre-training database consists of 172 neurons (77 neurons from monkey B, left hemisphere; 95 neurons from monkey Y, right hemisphere). As mentioned previously, a hexagonal stimulus array was used in the early pre-training experiments, and a rectangular array in the later pre-training experiments. Excluding the central location (which was reserved for the fixation light), there were 18 possible stimulus locations in the hexagonal array, and 24 possible stimulus locations in the rectangular array. Of the 172 cells in the pre-training database, 56 neurons (all from monkey B) were recorded while the hexagonal array was in use; data from the remaining 116 cells (21 neurons from monkey B, 95 neurons from monkey Y) were collected with rectangular stimulus array.

Post-training

The post-training database contains 160 neurons (99 neurons from monkey B, left hemisphere; 61 neurons from monkey Y, right hemisphere), for which all data were collected using two stimulus locations in the rectangular array.

2.2.8 Data analysis

Results of all analyses are pooled across the two monkeys, since all trends existed and reached significance in each monkey individually. Throughout the text, all statistical tests are two-tailed, and the critical significance level is 0.05 (n.s. means “not significant at the 0.05 significance level”). Analyses of neural data are based on firing rates over two intervals: the *pre-stimulus* period (the 500 ms interval before stimulus onset) and the *stimulus* period (the 500 ms interval from stimulus onset to stimulus offset). A neuron has a *significant spatially tuned response* if the firing rate in the stimulus period varies significantly across the stimulus locations tested (Kruskal-Wallis test). Analyses are designed to detect spatially tuned responses in the stimulus period, because spatially untuned responses cannot be distinguished from general arousal effects. However, analyses comparing firing rates in the stimulus and pre-stimulus periods at each location yielded similar results.

To estimate the power of the non-parametric Kruskal-Wallis test in the pre- and post-training experiments, a power analysis was performed separately for each experiment. First, the coefficient of variation in mean firing rate across stimulus locations was computed for each significant spatially tuned visual response. The average coefficient of variation ($\langle CV \rangle$) across all significantly tuned visual responses was taken as an estimate of the detectable dispersion in the data, and used to generate data sets with simulated spatial tuning. A simulated data set was created for each cell by randomly shifting the actual firing rates from auditory trials:

$$x_{nt}^{sim} = x_{nt} - \langle x_n \rangle + \langle x_n^{sim} \rangle \quad (2.1)$$

where x_{nt} and x_{nt}^{sim} are actual and simulated firing rates for trial t at location n , and $\langle x_n \rangle = \frac{1}{T} \sum_{t=1}^T x_{nt}$ is the actual mean firing rate for location n . The quantity $\langle x_n^{sim} \rangle$, representing the simulated mean firing rate for location n , is generated randomly from a normal distribution with mean $\frac{1}{N} \sum_{n=1}^N \langle x_n \rangle$ and standard deviation $\langle CV \rangle \frac{1}{N} \sum_{n=1}^N \langle x_n \rangle$. Thus artificial data sets containing simulated spatially tuned auditory responses were designed to have as much structure as was present in the average spatially tuned visual response. The proportion of simulated responses judged to have significant spatial tuning in the Kruskal-Wallis test was then taken to be a measure of the power of the test. The simulations and power calculations were repeated 100 times for both pre-training and post-training data sets.

An information theoretic analysis was used to quantify the degree of spatial tuning in stimulus-period firing rates for auditory or visual trials. Unlike the Kruskal-Wallis test, which simply categorizes a response as significantly tuned or not, the information theoretic analysis produces a continuous measure of spatial tuning. The distribution of this quantity, hereafter the *stimulus location information*, can then be used to summarize trends across the population of recorded neurons.

Estimates of stimulus location information were obtained as follows. For each set of auditory or visual trials collected from the same neuron, firing rates recorded at each stimulus location were used to build a matrix in which stimulus location constituted one dimension, and stimulus-period firing rate bins the other dimension. This matrix was then normalized to estimate joint probability and marginal probability densities. The width of the firing rate bins was defined to be the standard deviation of the pooled pre-stimulus-period firing rates. This quantity is an estimate of the noise in the cell's firing rate, and was used as the binwidth in order to obtain a conservative measure of the information content of the response (Gnadt and Breznen 1996). The mutual information between stimulus location and firing rate is given by:

$$I = \sum_s \sum_r P(s, r) \log_2 \frac{P(s, r)}{P(s)P(r)}, \quad (2.2)$$

where s is the index of each stimulus location, r is the index of the firing rate bin,

$P(s, r)$ is the joint probability, and $P(s)$ and $P(r)$ are the marginal probabilities (Cover and Thomas 1991).

Direct comparison of the mutual information between stimulus location and firing rate for LIP responses recorded before and after training is not possible, since there was a greater amount of location information available in the pre-training stimulus array than in the post-training array (due to the fact that 18–24 stimulus locations were used in pre-training experiments, but only 2 stimulus locations were used in post-training experiments). The stimulus location information must therefore be defined as the mutual information relative to some reference level of information. A conservative choice for this reference level is the median mutual information obtained after shuffling all the trials to destroy any correspondence between stimulus location and firing rate (Tovée et al. 1993). The shuffling procedure was repeated 100 times for each set of auditory or visual trials to estimate the median mutual information for the shuffled dataset. This reference level of information was then subtracted from the original mutual information (for unshuffled trials), to obtain the stimulus location information. Because the same data are used to calculate the original mutual information and the median trial-shuffled mutual information, this shuffle-subtracted measure of stimulus location information would not be subject to overestimation even if the firing rate bins were too small. Moreover, this quantity is appropriate for use in a hypothesis test to assess the degree of spatial tuning in firing rates across the recorded population. If the distribution of the stimulus location information across the population has a mean significantly greater than zero (Wilcoxon signed-rank test), then the firing rates carry significant information about stimulus location.

2.2.9 Histology

After the recordings in monkey B were completed, two electrolytic lesions, one shallow (2 mm below dura) and one deep (10 mm below dura), were placed in each of two electrode tracks. On the next day, monkey B was euthanized with pentobarbital sodium (50 mg/kg) and then perfused with a 10% formaldehyde solution. The brain

was extracted and sliced into 50 μm sections using a freezing microtome. Sections were then Nissl-stained with cresyl violet. Histological examination of the electrolytic lesion sites confirmed that recordings had been made in the lateral bank of the intraparietal sulcus. Figure 2.2 shows a photomicrograph of a section which includes one of the electrode tracks and lesion sites. Note that only the deep lesion is visible. This lesion is 10 mm below the dural surface, and at least 2 mm below the deepest recording sites.

2.3 Results

2.3.1 Pre-training

In the pre-training experiments, very few cells showed any response to auditory stimulation, but neurons with spatially tuned visual responses were often encountered. Figure 2.3 displays data from a cell typical of those recorded in area LIP before auditory-saccade training. The location of each plot in panels *A* and *B* corresponds to the location of a stimulus in the rectangular array. The plot in the middle of each grid (corresponding to the fixation point) displays representative fixational eye movements recorded during trials in which a stimulus was presented at location $(0^\circ, +16^\circ)$; note that the scale bar is approximately the width of the fixation window. Plots at all other locations in each panel show neural activity in response to a stimulus at the corresponding location on the stimulus array. Histograms indicate the firing rate of the neuron in spikes per second as a function of time relative to stimulus onset; the two bold vertical lines in each plot bracket the stimulus presentation interval. This cell has no auditory response at any location in the array (Figure 2.3 *A*), and no significant auditory spatial tuning (Kruskal-Wallis test, $p > 0.3$). However, the cell does have a strong response to visual stimulation (Figure 2.3 *B*), and this response is spatially tuned (Kruskal-Wallis test, $p < 0.001$).

Before auditory-saccade training, only 6 of 172 neurons, 3% of the cells recorded in area LIP, were found to have significant spatially tuned responses to auditory stimuli.

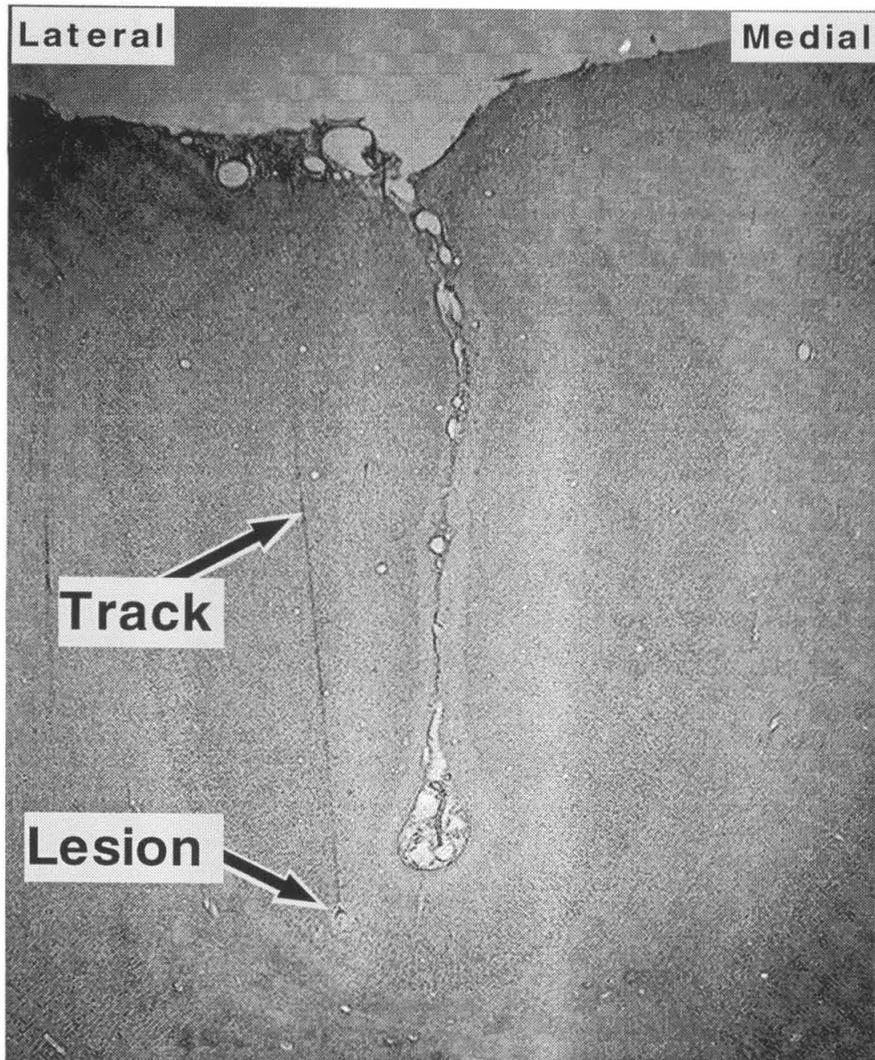


Figure 2.2: A Nissl-stained section of the intraparietal sulcus from the left hemisphere of monkey B. Left is the lateral aspect. One of the electrode tracks in which an electrolytic lesion was made is visible on the lateral bank of the sulcus. The lesion at the end of the electrode track is 10 mm below the dural surface, and at least 2 mm below the deepest recording sites.

PRE-TRAINING

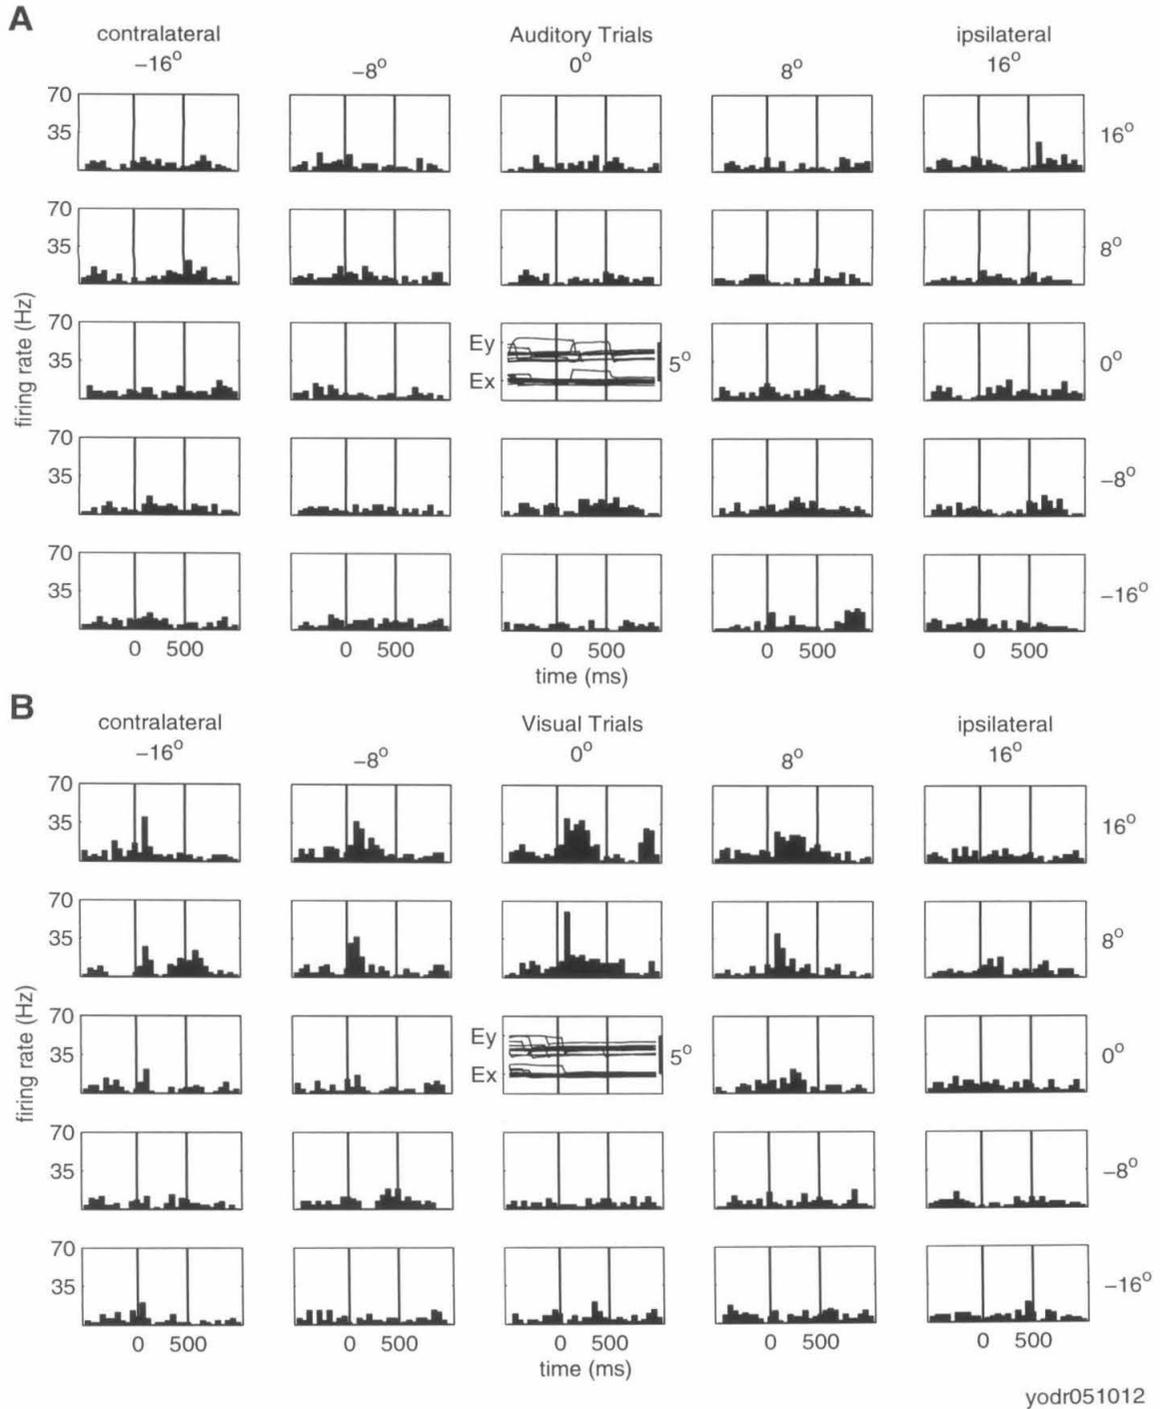


Figure 2.3: Activity of an LIP neuron typical of those recorded before auditory-saccade training. The cell has no evident response to auditory stimuli (A), but has a spatially tuned response to visual stimuli (B). See text for explanation of plots.

Given that the expected false positive rate for the Kruskal-Wallis test is 5%, this proportion is not significantly different from that expected by chance (binomial test, $p > 0.5$). In contrast, nearly half the cells (45%; 78/172) had spatially tuned visual responses, usually contralateral to the recording chamber. Percentages of neurons with spatially tuned responses are shown in Figure 2.4 for both auditory and visual stimulus modalities.

Visual inspection of data from the 6 cells judged to have significant spatially tuned auditory responses confirmed that even these cells did not have obvious responses to auditory stimuli. Figure 2.5 displays the activity of the cell with the smallest p -value in the test for spatially tuned auditory responses (Kruskal-Wallis test, $p = 0.0008$). There is no discernible response to auditory stimulation (Figure 2.5 *A*).

To quantify the spatial selectivity of auditory and visual responses across the recorded population, an information theoretic analysis was performed (see Methods). For each cell, the mutual information between stimulus location and stimulus-period firing rate was compared to the median mutual information obtained after trials had been shuffled to destroy any correspondence between stimulus location and firing rate. As explained in the Methods section, the difference between these two mutual information quantities (the stimulus location information) is a conservative estimate of the amount of location information carried in the stimulus-period firing rates. The distribution of this information measure across all cells in the pre-training database is shown for auditory and visual trials in Figure 2.6 *A* and *B*. The mean of the distribution for auditory trials (Figure 2.6 *A*) is not significantly different from zero (Wilcoxon signed rank test, $p > 0.6$), indicating that firing rates in the pre-training database do not carry significant amounts of information about auditory stimulus location. In contrast, the mean of the distribution for visual trials is significantly greater than zero ($p < 0.001$). Therefore, the firing rates of LIP neurons recorded before auditory-saccade training do not convey spatial information regarding auditory stimuli, but do carry information about the locations of visual stimuli.

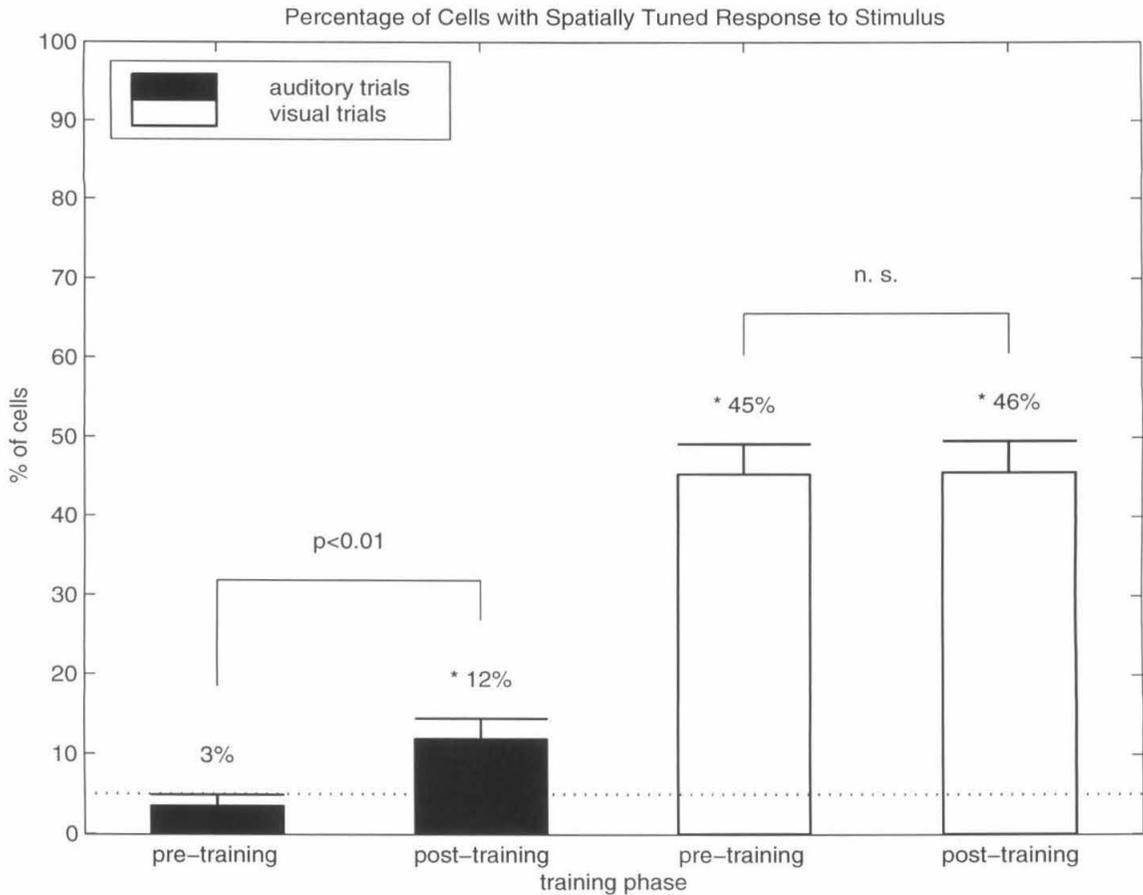


Figure 2.4: Percentages of recorded cells with spatially tuned responses to auditory or visual stimuli, in the pre-training and post-training experiments. The dotted line indicates the expected false positive rate for the Kruskal-Wallis test used to identify cells with spatially tuned responses, and asterisks mark proportions which are significantly different from this chance level (binomial test). Probabilities above brackets show the significance level for the difference between adjacent proportions (Fisher-Irwin test). Error bars represent theoretical standard deviation, estimated from the observed response percentage. Total number of cells: pre-training $N = 172$, post-training $N = 160$.

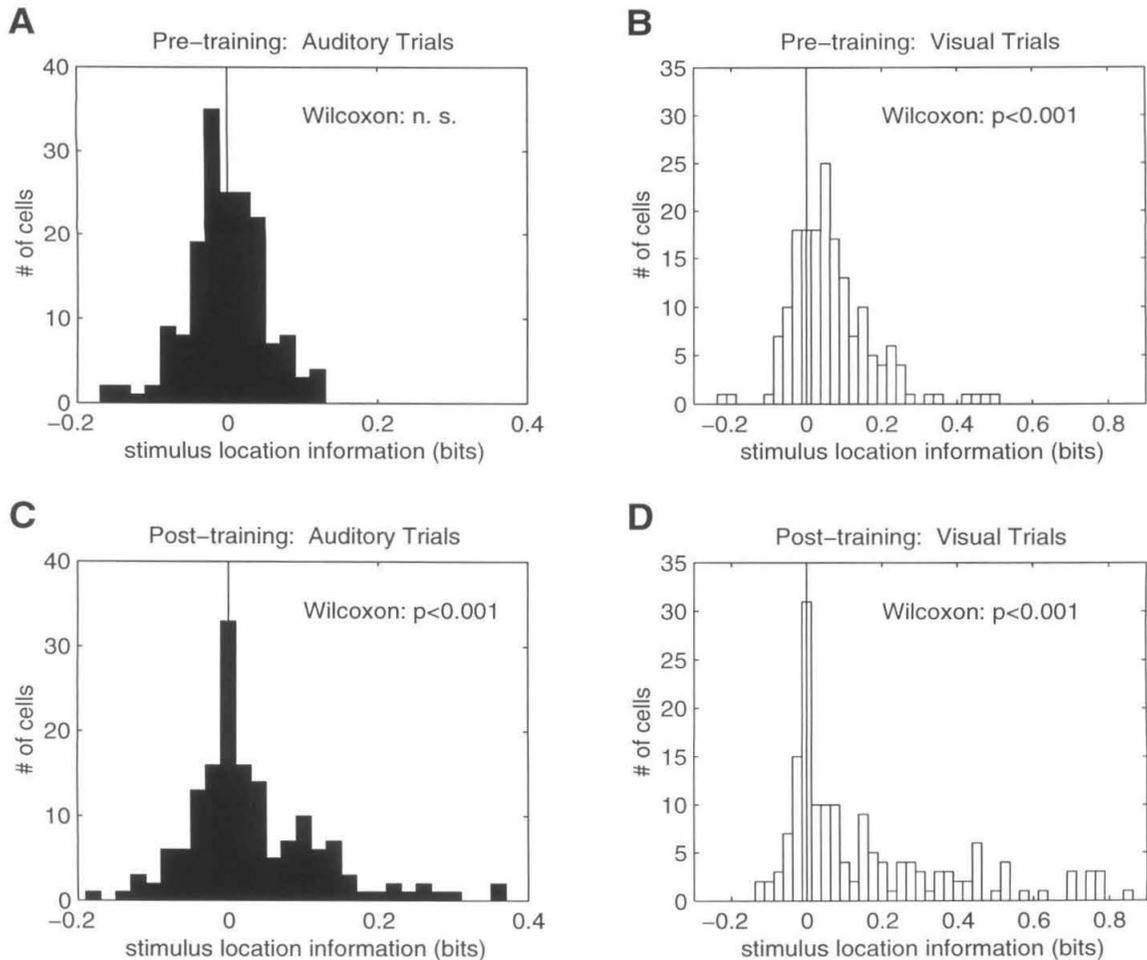


Figure 2.6: Location information available in stimulus-period firing rates, for all cells in the pre- and post-training databases. The vertical line in each plot marks zero difference between original trials and shuffled trials (see Methods). Wilcoxon signed-rank test results indicate the significance level for rejection of the hypothesis that the mean of the distribution is zero (no stimulus location information across the population). Firing rates recorded during auditory trials carry significant stimulus location information only in the post-training database (*C*), not in the pre-training database (*A*). In contrast, firing rates during visual trials convey significant spatial information both before and after auditory-saccade training (*B* and *D*). Total number of cells: pre-training $N = 172$, post-training $N = 160$.

2.3.2 Post-training

In the post-training experiments, some neurons exhibited very brisk responses during auditory stimulation, while the monkey was performing the modified fixation task. Figure 2.7 shows such a neuron. In addition, many cells had visual responses, usually contralateral to the recording chamber.

Across the population, about 12% (19/160) of neurons recorded after auditory-saccade training had significant spatially tuned responses to auditory stimulation. Half the cells (46%; 73/160) exhibited spatially tuned responses to visual stimulation. These proportions are illustrated in Figure 2.4. The percentage of cells with spatially tuned auditory responses is significantly greater than the expected false positive level, indicated by the dotted line (binomial test, $p < 0.001$).

Stimulus location information estimates for the post-training database are shown in Figure 2.6. The mean of the population distribution for auditory trials is significantly different from zero (Wilcoxon signed rank test, $p < 0.001$), indicating that firing rates during auditory trials carry significant information about auditory stimulus location after auditory-saccade training. For visual trials, the mean of the distribution is also significantly greater than zero ($p < 0.001$). Thus, in the post-training database, firing rates during both auditory and visual trials convey significant spatial information.

2.3.3 Comparison between pre- and post-training

As indicated in Figure 2.4, the proportion of recorded neurons with significant spatially tuned responses to auditory stimuli is significantly higher after auditory-saccade training than before (Fisher-Irwin test, $p < 0.01$). However, there is no significant difference between the proportions of cells with spatially tuned visual responses before and after auditory-saccade training ($p > 0.05$).

A direct comparison of the stimulus location information available in neural firing rates before and after training is not possible, because more stimulus locations were used in the pre-training experiments than in the post-training experiments. However,

POST-TRAINING

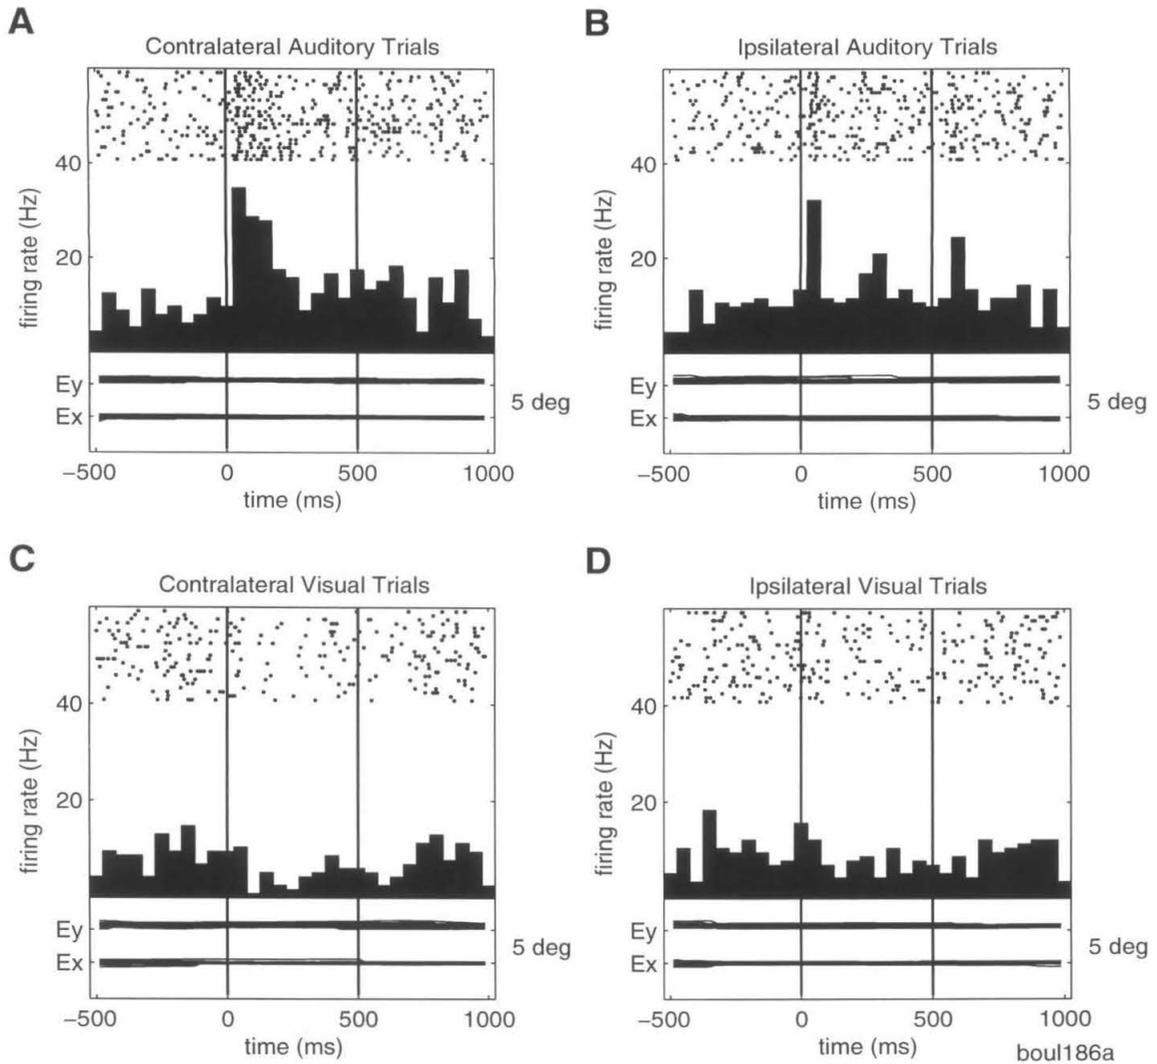


Figure 2.7: Activity of an LIP neuron recorded after auditory-saccade training. Rasters at the top of each plot indicate the times of spike occurrence in each trial; histograms show the firing rate of the neuron as a function of time relative to stimulus onset; and eye position traces indicate horizontal (E_x) and vertical (E_y) eye position during each trial. Stimulus locations are specified as contralateral (A and C) or ipsilateral (B and D) to the recording chamber. This neuron has a significant spatially tuned response to auditory stimuli (Kruskal-Wallis test, $p < 0.05$). Visual response tuning does not reach significance, although some inhibition is evident during contralateral visual trials (C).

it is appropriate to ask if there is a significant amount of stimulus location information available in neural firing rates within each database. As explained previously, firing rates in area LIP convey no information about auditory stimulus location before auditory-saccade training (Figure 2.6 *A*; Wilcoxon signed-rank test, $p > 0.6$), but do carry significant information about visual stimulus location (Figure 2.6 *B*; $p < 0.001$). After auditory-saccade training, firing rates during both auditory and visual trials carry significant amounts of stimulus location information (Figure 2.6 *C* and *D*; both trial types, $p < 0.001$). Thus the information theoretic analysis confirms that the firing rates of neurons in area LIP carry no information about the locations of auditory stimuli before auditory-saccade training, but that auditory spatial information is represented in area LIP after training.

2.3.4 Control for search/selection bias

Recording strategies before and after training were slightly different, as explained in the Methods section. It is possible that differences in the search tasks and cell selection criteria could account for the apparent increase in the auditory responsiveness of LIP after auditory-saccade training. Two additional analyses argue against this possibility.

First, the differences between the pre- and post-training databases persist even when an artificial *post hoc* search bias is introduced, restricting the analysis to visually responsive cells. In this restricted analysis, only neurons with spatially tuned visual responses are considered, and the proportions of cells with spatially tuned responses to auditory stimuli are calculated for this subsample. Among the cells recorded before auditory-saccade training, 78 had spatially tuned visual responses, and of these only 2 responded to auditory stimuli. In contrast, of the 73 cells with spatially tuned visual responses after auditory-saccade training, 13 also had spatially tuned responses to auditory stimuli. These proportions are significantly different (2/78 vs. 13/73; Fisher-Irwin test, $p < 0.005$), indicating that in this limited sample there are more responses to auditory stimuli after training than before. While not conclusive, this analysis suggests that across the population, the number of cells with spatially tuned

auditory responses increased due to auditory-saccade training.

Second, a control experiment demonstrates that recording strategies alone cannot account for the results. An additional 33 cells were recorded in monkey Y after auditory-saccade training, using the same recording strategy employed in the pre-training experiments. In this control experiment, the search task was the fixation task, and data were collected from every cell that was isolated. Among all the neurons recorded in this way, 12% (4/33) had significant spatially tuned responses to auditory stimuli, a significantly higher percentage than would be expected by chance (binomial test, $p < 0.05$). Thus differences in the recording strategies are an unlikely explanation for the increased number of responses to auditory stimuli found after auditory-saccade training.

2.3.5 Control for different arrays

Before training two different stimulus arrays were employed, one hexagonal and one rectangular (see Methods). The majority of cells in monkey B were recorded using the hexagonal array, while all of the cells in monkey Y were recorded using the rectangular array. The proportion of cells with auditory spatial tuning was about the same for both arrays (3%), and no more than expected by chance for either array (binomial test, $p > 0.2$). The proportion of cells with spatially tuned visual responses was 59% for the hexagonal array and 37% for the rectangular array.

As noted in the Methods section, the frequency spectra of the speakers were matched in the post-training experiments but not in the pre-training experiments. It is unlikely that this difference contributed to the increase in LIP responsiveness in the post-training experiments, because matching made the auditory stimuli more uniform across the array. If anything, matching should have reduced, not increased, variation in auditory responses across locations, and hence should have reduced the apparent spatial tuning of auditory responses in area LIP.

2.3.6 Control for number of locations/samples

In the pre-training experiments, the number of stimulus locations was considerably higher than in the post-training experiments. To equalize the two datasets, the pre-training database was restricted to the two locations used in the post-training experiments (or the two closest locations for cells recorded using the hexagonal array). The spatial tuning analysis was then repeated using only those two locations from the pre-training database, allowing direct comparison of pre- and post-training auditory responses. Only 2% of cells in the restricted pre-training dataset exhibited spatially tuned responses to auditory stimuli. Comparison between the restricted pre-training dataset and the post-training data reveals a significant increase in the proportion of cells with spatially tuned responses to auditory stimuli (Fisher-Irwin test, $p < 0.001$). Thus it is unlikely that the apparent effect of training is an artifact of spatial undersampling in the post-training experiments.

Because many more stimulus locations were used in the pre-training experiments than in the post-training experiments, the number of repetitions per location tended to be lower pre-training (between 5 and 10) than post-training (between 10 and 20). The power of a test is increased both by the number of conditions, and by the number of samples. Since there were more locations and fewer samples per location in pre-training experiments, it is conceivable that the power of the Kruskal-Wallis test might have been lower pre-training than post-training. Such a difference in power would make responsive cells less likely to be detected in pre-training experiments. This scenario seems unlikely, since any power difference should have affected the visual responses too, and the proportion of neurons with significant spatially tuned visual responses was about the same pre- and post-training. However, to address this issue more directly, the power of the spatial tuning test was estimated in Monte Carlo simulations (see Methods section for details). Over 100 simulations, the average power of the Kruskal-Wallis test to detect comparable differences was 77% (range: 72–81%) before training, and 46% (range: 31–54%) after training, indicating that the larger number of locations used in the pre-training experiments outweighs the

smaller number of repetitions. Thus the absence of responses to auditory stimuli before training is not due to lower statistical power in the test for spatial tuning before training. In fact, this power analysis suggests that responses to auditory stimuli were less likely to be detected in the post-training experiments; in other words, the estimated proportion of cells responding to auditory stimuli in the post-training experiments likely underestimates the true proportion.

2.3.7 Control for post-reward eye movements

Given that responses to auditory stimuli are more prevalent in a saccade task than in a fixation task (see Chapter 3), it is possible that responses to auditory stimuli might appear in fixation trials if the monkeys were performing very late saccades to the stimulus locations in the post-training experiments. If this were the case, then the apparent increase in auditory responsiveness of LIP after training might be due to eye movements, not due to training. To address this concern, eye position was recorded for at least 500 ms after the reward, without any behavioral constraint on the monkey. Saccades during this period were detected using velocity criteria, and the eye position after the first saccade was extracted. If no saccade occurred, the eye position at the end of the recording period was used. Figure 2.8 shows these post-reward eye positions for fixation trials collected from the cell shown in Figure 2.7. After the reward, the monkey did not continue fixating, but tended to make an eye movement up and to the right, presumably toward his default eye position. The final eye position distributions are similar for trials in which the stimulus appeared on the left and trials in which the stimulus appeared on the right. A similar analysis was performed for all recordings from neurons with significant spatially tuned auditory responses in the post-training experiments. For not a single neuron with auditory spatial tuning did post-reward eye positions differ depending on the location of the stimulus (Kolmogorov-Smirnov test separately for the horizontal and the vertical dimensions). In other words, late, goal-directed saccades cannot explain the increased auditory responsiveness of area LIP after auditory-saccade training.

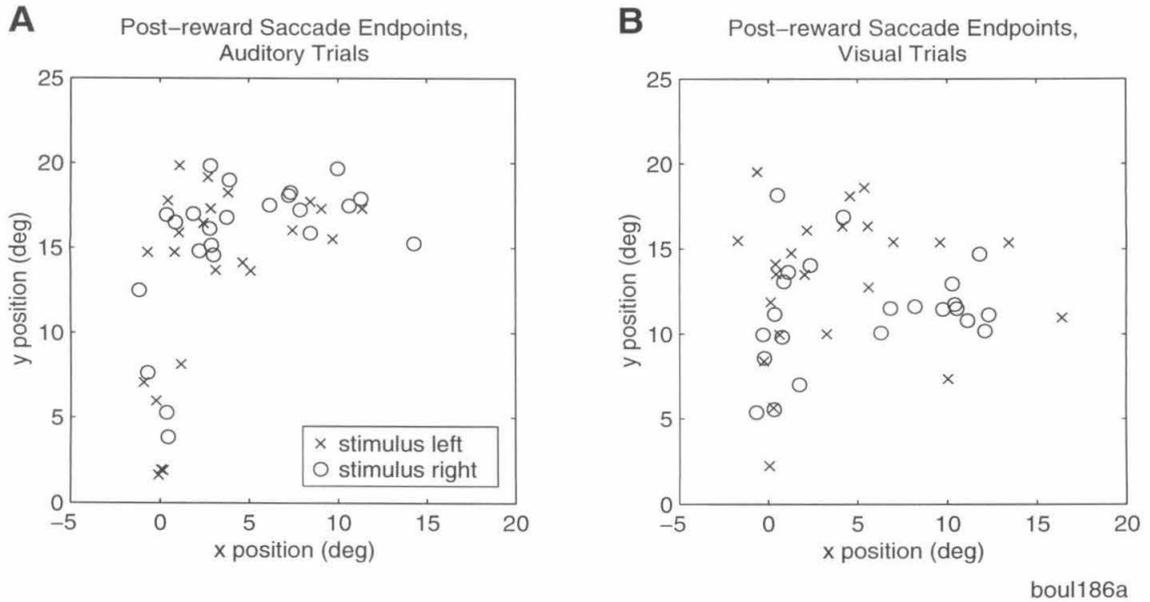


Figure 2.8: Eye position after the first post-reward saccade, for trials collected in the same post-training recording session as the data shown in Figure 2.7. Crosses indicate trials in which the stimulus was to the left of fixation, circles indicate trials in which the stimulus was to the right of fixation. The difference between the two distributions did not reach significance for either auditory (*A*) or visual (*B*) trials (one-dimensional Kolmogorov-Smirnov test along the horizontal dimension, $p > 0.5$ for both auditory and visual trials).

2.3.8 Control for penetration locations

To check that the recordings were made from approximately the same locations before and after training, penetration maps were constructed. Figure 2.9 shows the types of responses that were associated with each recording site. Most recording sites were penetrated at least ten times in the pre-training experiment. As a result, especially in monkey B, many penetration sites in the center of the chamber that had been visually responsive in the pre-training experiments did not respond visually in the post-training experiments, probably due to tissue damage. Nevertheless, the pre- and post-training maps largely overlap, and the locations of most responses suggest that the recordings stem from the same brain location. The map also indicates the approximate location of the brain section shown in Figure 2.2, and the location of the electrolytic lesion visible in that section.

2.4 Discussion

2.4.1 Effects of auditory-saccade training

The key finding of this chapter is that auditory-saccade training increases the responsiveness of LIP to auditory stimuli. This result is demonstrated in two independent analyses. In the first analysis, each response is categorized as exhibiting significant or no significant spatial tuning, and the proportions of neurons with significant spatially tuned responses before and after training are compared. In the second method, the amount of information about stimulus location available in firing rates across the population is estimated, both before and after training. The two methods obtain their results differently, but arrive at the same conclusion. There are, however, several possible confounds that must be addressed, before the effect of training can be viewed as established.

First, it is possible that the search task, and the criteria by which neurons were selected for further recording, biased the results in such a way as to inflate the number of cells that exhibited spatially tuned responses to auditory stimuli after training. As

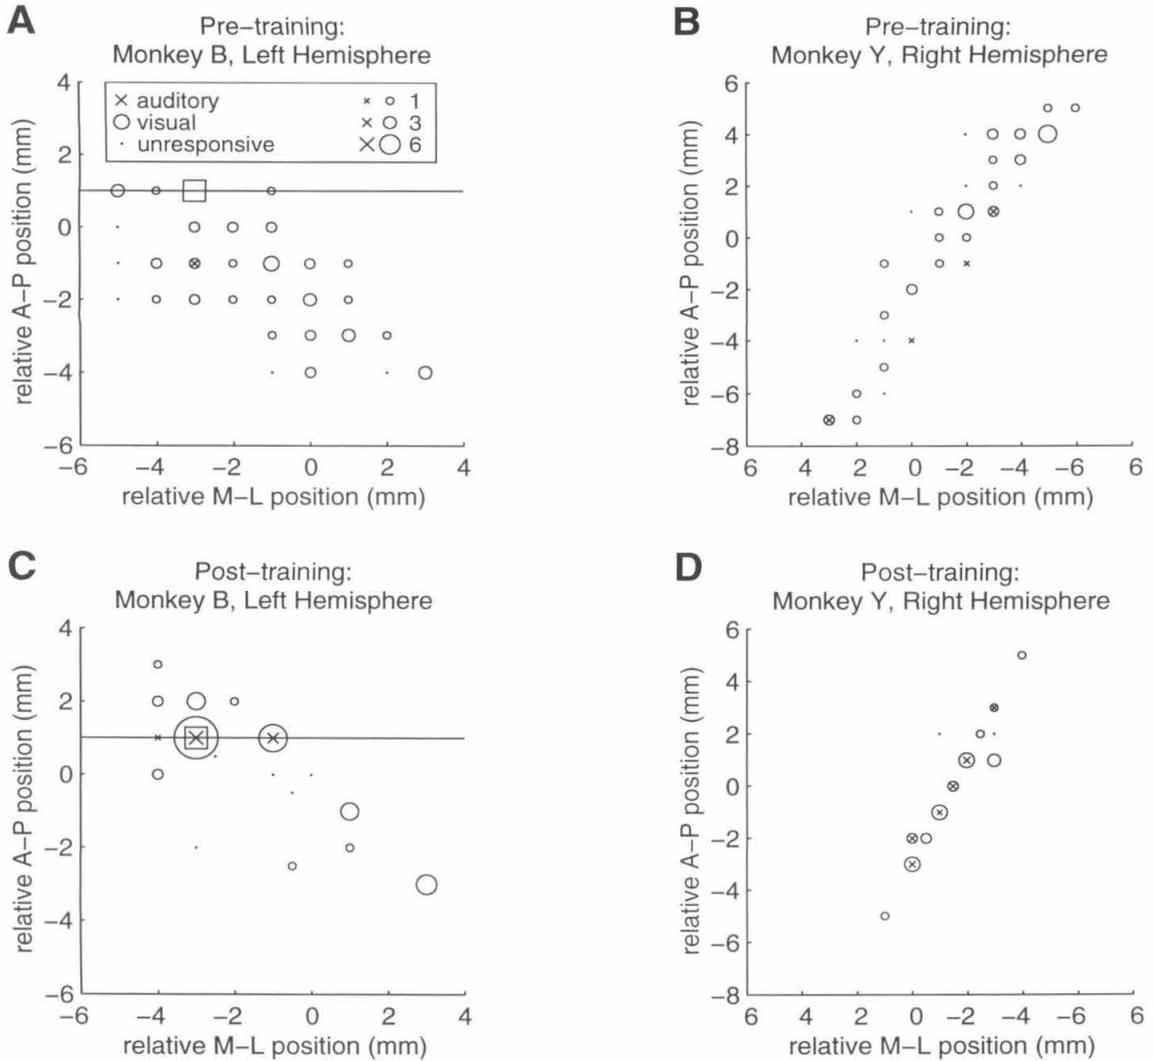


Figure 2.9: Distributions of neurons with significant spatially tuned auditory or visual responses, in the pre- and post-training experiments. Each plot shows the positions of all electrode penetration sites, in mm anterior–posterior (A-P) and mm medial–lateral (M-L) relative to the center of the recording chamber. Symbols at each penetration site indicate the type of spatially tuned responses observed: auditory (\times), visual (\circ), or neither (\cdot). The sizes of cross and circle symbols at each site are scaled to reflect the number of neurons with the corresponding type of response. *A* and *C*: Pre- and post-training electrode penetration sites for monkey B, whose recording chamber was mounted over the left hemisphere. The square shows the site of one of the electrolytic lesions made in this animal, and the line indicates the approximate angle of the histological section shown in Figure 2.2. *B* and *D*: Pre- and post-training penetration sites for monkey Y, whose recording chamber was mounted over the right hemisphere.

discussed in the Results section, two lines of reasoning, including a control experiment, argue against this interpretation.

Second, the apparent effect of training might be due to the limitation to a smaller stimulus array in the post-training task. Since this difference would have made responses more difficult to detect in the post-training phase, and since the power of the Kruskal-Wallis test was actually higher before training, this is an unlikely explanation for the training effect.

Third, in light of the data presented in Chapter 3, the task that monkeys are performing appears to have a strong impact on the auditory responsiveness of LIP neurons. In the post-training experiments, blocks of trials in which the monkeys performed the modified fixation task were alternated with blocks in which they performed memory saccades; therefore, it is possible that the animals were making saccades after receiving a reward in the fixation task. However, as shown in the Results section, final eye positions in the fixation task following the first eye movements after the reward did not differ depending on the stimulus location. Thus goal-directed, post-reward eye movements are an unlikely explanation for the post-training responsiveness of LIP to auditory stimuli.

Finally, it is possible that the pre- and post-training recordings were made from different areas, and that in fact the comparison is not valid. This is an issue because the pre-training and post-training experiments spanned about two years, during which brain and skull growth might have affected the position of cortical areas. In as much as the brain remains more or less at the same location relative to the recording chamber over an extended period of time, the penetration maps in the Results section show that the pre- and post-training recording sites overlap substantially.

In summary, the results and controls indicate that the responsiveness of LIP neurons to auditory stimuli in a fixation task changed as a consequence of training the animals to perform auditory saccades. Thus one explanation for the apparent discrepancy between early studies that reported no auditory responses in LIP (Mountcastle et al. 1975; Hyvärinen 1982b; Koch and Fuster 1989) and more recent studies that did find auditory responses in LIP (Mazzoni et al. 1996b; Stricanne et al. 1996) is

that in the latter studies the monkeys had been trained to perform auditory saccades, while in the former studies they had not.

In the introduction to this chapter, four possible scenarios were proposed in which the discrepancies between the early and later studies of auditory responses in area LIP might be resolved. The first possibility was that auditory responses exist in area LIP before auditory-saccade training, but had been overlooked in the earlier studies. The second possibility was that auditory-saccade training induces responses to auditory stimuli in area LIP. The third possibility was that the task an animal is performing determines the auditory responsiveness of area LIP. Finally, the fourth possibility was that both training and task affect neural responses to auditory stimuli in LIP. The present results support the second and fourth possibilities. In light of the results of the next chapter, which show that responses to auditory stimuli in area LIP are modulated by behavioral task, the fourth possibility — that responses to auditory stimuli in LIP are affected both by training and by the task the animal is performing — appears to be correct.

2.4.2 Possible mechanisms

At the mechanistic level, the observed effect of training may arise through emergence of new connections to area LIP from an auditory area that is as yet unidentified. Alternatively, it is possible that training unmasks connections that existed all along but were silent. In both cases, several different areas may be providing auditory input to area LIP. Likely candidate regions in cortex are temporoparietal cortex (area Tpt), the superior temporal polysensory area (STP), and the frontal eye fields (FEF). A possible subcortical source of auditory input to LIP is the deep layers of the superior colliculus. All of these regions respond to auditory stimuli (Leinonen et al. 1980; Russo and Bruce 1994; Jay and Sparks 1984; Hikosaka et al. 1988). Areas Tpt, STP, and FEF project directly to area 7, including area LIP (Blatt et al. 1990; Pandya and Kuypers 1969; Baizer et al. 1991), while the superior colliculus projects to LIP via the pulvinar (Asanuma et al. 1985).

Neurons in areas Tpt and STP respond to auditory stimuli even in animals which have not been trained to perform auditory saccades (Leinonen et al. 1980; Baylis et al. 1987), and both regions receive strong projections from other auditory areas (Kaas and Hackett 1998). Frontal eye field neurons respond to auditory stimuli in a fixation task after auditory-saccade training (Vaadia et al. 1986), but it is unclear whether or not FEF neurons respond to auditory stimuli before auditory-saccade training, or in anesthetized animals. In light of the present results in area LIP, it seems possible that auditory responses in FEF also emerge through training, especially since there are strong projections from LIP to FEF (Andersen et al. 1990; Asanuma et al. 1985). Assuming that the areas providing direct or indirect auditory input to LIP respond to auditory stimuli without behavioral training, areas Tpt and STP seem the most likely sources of auditory input to LIP, although the auditory signals might well be routed through frontal cortex. It is also possible that the superior colliculus could be the source of the responses to auditory stimuli in LIP, since the deep layers of the colliculus respond to auditory stimuli in anesthetized monkeys (Cynader and Berman 1972).

2.4.3 Relation to previous work

The present study is the first report demonstrating emergence of responses to auditory stimuli *de novo* in posterior parietal cortex after saccade training. Similar training-induced increases in responsiveness have been reported in area 3a following tactile discrimination training (Recanzone et al. 1992). In addition, visual search training using color cues has been shown to induce color selectivity in neurons of the frontal eye fields (Bichot et al. 1996).

The effects of auditory-saccade training in this study are reminiscent of training-induced changes observed in frontal cortex by Watanabe (1992), who showed that prefrontal neurons code the associative significance of auditory and visual stimuli. In that study, a cue stimulus indicated to the animal whether a subsequent trial would be a rewarded or an unrewarded trial. After training, auditory or visual cues signalling

a rewarded trial evoked stronger firing in prefrontal neurons than cues signalling an unrewarded trial. Clearly, the cues were behaviorally relevant to the animal, but the cue stimuli informed the animal only of the outcome of the subsequent trial. The training effect observed in the present study is somewhat different, in that it occurred in the context of eye movements, a context which is likely to be critical to LIP function.

In a different study, Chen and Wise (1995a) showed that neurons in the supplementary eye fields (and, to a lesser extent, the frontal eye fields; Chen and Wise 1995b) code conditional oculomotor associations between random visual stimuli and upcoming eye movements. Learning occurred within one session in that study, and thus the neural activity could be studied at the same time as learning took place. Neurons which initially had not responded to novel visual stimuli indicating the direction of an upcoming saccade started responding during training (Chen and Wise 1995a), and developed selectivity for the upcoming saccade as the session progressed (Chen and Wise 1996). The effects of training observed in the present experiments seem to be very similar to those observed in frontal cortex by Chen and Wise. However, in the present experiments, training occurred over a time span of several months, while Chen and Wise trained their animals during single recording sessions lasting at most several hours. Further research will be necessary to determine the exact correspondence between long-term training-induced changes in area LIP, and the short-term conditional oculomotor associations observed in frontal cortex.

2.4.4 Interpretations

The dependence of auditory responses in area LIP on auditory-saccade training suggests that these responses cannot really be termed "sensory auditory" responses. Rather, the responses are contingent upon the monkey being trained in an auditory-saccade task. How might responses to auditory stimuli emerge through training? Four cognitive-level explanations will be considered.

First, the emergence of responses to auditory stimuli may reflect a change in the

attentional state of the animals. It is possible that the monkeys were completely ignoring the auditory stimuli before training. After auditory-saccade training, when they had learned that the auditory stimuli could be behaviorally relevant, the animals may have paid attention to the stimuli in the fixation task even though they were not required to do so. Indeed, other investigators have suggested that LIP signals reflect the allocation of attentional resources (Colby et al. 1996; Gottlieb et al. 1998).

Alternatively, a change of intention *vis à vis* auditory stimuli may underlie the effects of auditory-saccade training on area LIP. Before training, the auditory stimuli were presumably irrelevant to the animals as saccade targets; through training, they became associated with saccades. Responses to auditory stimuli in LIP may reflect covert plans to make eye movements to auditory stimuli, plans which are formed even when the animal is instructed not to make any eye movements. Previous studies have demonstrated that a component of LIP activity signals the intention to make eye movements (Bracewell et al. 1996; Mazzoni et al. 1996a), and that in the absence of actual movements, activity in LIP may code the intention but not the execution of eye movements (Snyder et al. 1997). In fact, the eye movement plan can be changed without any movement being executed, and activity in LIP reflects these changes (Snyder et al. 1998). Moreover, LIP activity in response to a visual stimulus quickly fades if the stimulus is identified as irrelevant in a saccade task (Platt and Glimcher 1997b; Shadlen and Newsome 1996).

A third possibility is situated between the attentional and intentional interpretations, and posits that auditory activity after saccade training codes the oculomotor significance of auditory stimuli — that is, the significance of auditory stimuli as potential saccade targets. In the present experiments, monkeys were trained that auditory stimuli had a meaning or significance in terms of oculomotor behavior. As a result, these stimuli may have become represented in area LIP. This idea is consistent with the observation that increased probability that a stimulus will be an eye movement target, or increased reward associated with a particular stimulus, strengthens the representation of that stimulus in area LIP (Platt and Glimcher 1997a). This interpretation could also explain the shape selectivity recently reported in area LIP of

animals trained to use shape stimuli as targets in an eye movement task (Sereno and Maunsell 1998). Moreover, the oculomotor significance idea could be extended to explain why LIP neurons are so responsive to visual stimuli even when eye movements are not required; easily localized visual stimuli might have default oculomotor significance. Similarly, some sounds (e.g., species-specific warning calls, or sounds from behind the animal) may have much higher oculomotor significance than the auditory stimuli used in the present study, and hence might elicit responses to auditory stimuli from LIP even before auditory-saccade training.

The fourth interpretation is that attempts to compartmentalize attention, intention and oculomotor significance are artificial. This fourth possibility would posit that sensory and movement activation should be expected to co-occur, and that it may not be useful to assign the post-training auditory activity in area LIP to any of the other three interpretations. Andersen and collaborators have previously argued that the parietal cortex participates in sensory-motor processing, operating as an interface between sensory and motor systems to transform sensation into action (Andersen et al. 1997). Attention likewise has been proposed to have evolved from circuits for orienting toward stimuli, and attentional mechanisms may well serve the purpose of preparing for action (Rizzolatti et al. 1994; Kustov and Robinson 1996). Indeed, eye movement and attention circuits are largely overlapping in the primate cortex (Corbetta et al. 1998). Interconnections and similarities in physiology between LIP and the frontal eye fields (Chafee and Goldman-Rakic 1998) suggest that sensory attention and movement planning share the same circuits, and therefore may not be modular and separate operations.

The present study cannot distinguish between these four interpretations of responses to auditory stimuli in LIP, because the experiments reported here were not aimed at distinguishing between them. Instead, the experiments were designed to examine why earlier reports did not find auditory responses in LIP, and more recent studies, using animals performing delayed auditory-saccade tasks, did. Future research will be needed to determine which of the four possibilities outlined above is the best interpretation of the emergence of responses to auditory stimuli in area LIP

after auditory-saccade training.

Chapter 3 Behavioral Modulation

The lateral intraparietal area was once thought to be unresponsive to auditory stimulation. However, recent studies have shown that neurons in area LIP respond to auditory stimuli during an auditory-saccade task. To what extent are responses to auditory stimuli in area LIP dependent on the performance of an auditory-saccade task? To address this question, recordings were made from 160 LIP neurons in two monkeys while the animals performed auditory and visual memory-saccade and fixation tasks. Responses to auditory stimuli were significantly stronger during the memory-saccade task than during the fixation task, while responses to visual stimuli were not. Neurons responsive to auditory stimuli tended also to be visually responsive, and to exhibit delay or saccade activity in the memory-saccade task. These results indicate that, in general, auditory responses in area LIP are modulated by behavioral context, are associated with visual responses, and are predictive of delay or saccade activity. Responses to auditory stimuli in area LIP may therefore be best interpreted as supramodal responses, and similar in nature to the delay activity, rather than as modality-specific sensory responses. The apparent link between auditory activity and oculomotor behavior suggests that the behavioral modulation of responses to auditory stimuli in area LIP reflects the selection of auditory stimuli as targets for eye movements.

3.1 Introduction

As explained in Chapter 2, early physiological investigations of LIP and surrounding regions found no auditory activity in this area (Mountcastle et al. 1975; Hyvärinen 1982b; Koch and Fuster 1989). However, more recent studies have shown that neurons in area LIP do respond to auditory stimuli in the context of an auditory-saccade task (Mazzoni et al. 1996b; Stricanne et al. 1996). There are several possible explanations for this apparent discrepancy. One possibility is that neurons which respond to auditory stimulation exist in area LIP, but were overlooked in early studies of posterior parietal cortex. A second possibility is that LIP neurons respond to auditory stimuli after auditory-saccade training, regardless of the immediate behavioral context of auditory stimulation after training. A third possibility is that neurons in area LIP respond to auditory stimuli only when the animal is engaged in an auditory-saccade task. Finally, a fourth possibility is that LIP neurons develop responses to auditory stimuli through auditory-saccade training, and subsequently display auditory activity primarily but not exclusively during an auditory-saccade task. Auditory responses of this type would be affected both by the animal's training history and by the immediate behavioral context in which an auditory stimulus appeared after training.

Chapter 2 excludes the first and third of these four possibilities, by demonstrating both that auditory responses do not appear in area LIP before auditory-saccade training, and that auditory responses are observed after training even when the animal is just fixating. The present chapter addresses the second and fourth possibilities, which concern the effects of immediate behavioral context on auditory responses in the trained animal. The experiments show that neurons in area LIP respond more strongly to auditory stimuli when monkeys are engaged in a memory-saccade task than when they are engaged in a fixation task. This behavioral modulation of auditory responses resembles behavioral modulation of delay-period activity. The experiments also reveal that LIP neurons with auditory responses tend to have visual responses, and to exhibit delay or saccade activity. Together, Chapters 2 and 3 demonstrate

that responses to auditory stimuli in LIP are dependent both on long-term training history and on short-term behavioral context. Furthermore, the results suggest that auditory responses in area LIP are best considered supramodal (cognitive or motor) responses, rather than modality-specific sensory responses. Task-dependent increases in responses to auditory stimuli in area LIP seem to reflect the selection of auditory stimuli as targets for eye movements. Reports of these results have been published in abstracts (Grunewald et al. 1997; Linden et al. 1998) and in an article (Linden et al. 1999).

3.2 Methods

3.2.1 Animals, animal care, and surgical procedures

Animals, animal care, and surgical procedures, explained in detail in Chapter 2, are summarized only briefly here. Two adult male *Macaca mulatta* monkeys were used as subjects in this study. A stainless steel head post, dental acrylic head cap, scleral search coil, and stainless steel recording chamber were implanted in each monkey using standard techniques (Mountcastle et al. 1975; Judge et al. 1980). The recording chamber was mounted normal to the surface of posterior parietal cortex (stereotaxic coordinates at center: 6 mm posterior, 12 mm lateral) over the left hemisphere of monkey B, and over the right hemisphere of monkey Y. After surgery, monkeys were given at least one week to recover before behavioral training or recording began. All surgical procedures and animal care protocols were approved by the California Institute of Technology Institutional Animal Care and Use Committee, and were in accordance with National Institutes of Health guidelines.

3.2.2 Experimental setup

The experimental setup is described in Chapter 2. Briefly, all experiments were conducted in complete darkness, in a double-walled sound-attenuating anechoic chamber (Industrial Acoustics Company, Inc.). While inside the chamber, the monkey was

monitored continuously with an infrared camera and a microphone. The animal faced a fixed stimulus array consisting of a concave rectangular grid of concentrically mounted piezoelectric speakers and light-emitting diodes. The monkey's head was held fixed during all behavioral training and recording sessions. Locations of stimuli are specified relative to the center of the monkey's head, in degrees azimuth right or left of the median sagittal plane and in degrees elevation above or below the visual plane. All stimuli in the concave stimulus array were approximately 80 cm from the monkey's head.

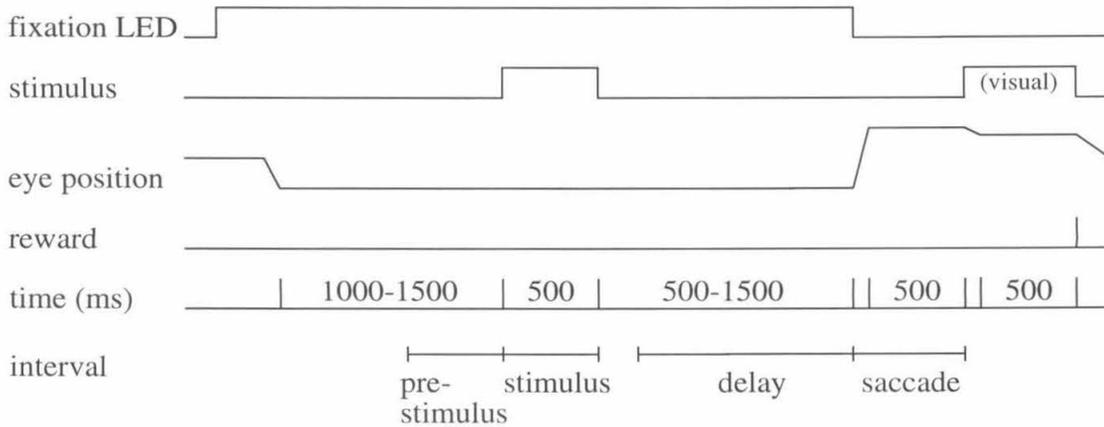
Visual stimuli were 500 ms flashes of 70 cd/m² red light from the LEDs, which each subtended 0.4 degrees. Free-field auditory stimuli were 500 ms bursts of band-limited noise (5–10 kHz, 5 ms rise/fall times, 70 dB SPL). This noise band was chosen because macaque monkeys have been reported to localize 5–10 kHz bandlimited noise well in azimuth (Brown et al. 1980), and because the frequency responses of the speakers were relatively flat (± 10 dB SPL) within this range. For most of the experiments reported here, the input to each speaker was adjusted to equalize the output amplitude spectrum to 70 ± 2 dB SPL within the 5–10 kHz frequency band, as measured at the location where the monkey's head would be during an experiment. There were no qualitative differences in behavioral or neurophysiological results obtained before and after the speakers were equalized.

3.2.3 Behavioral paradigms

Neural recordings were obtained while the monkeys were performing two tasks: the memory-saccade task and the fixation task (Figure 3.1). Two fixed stimulus locations were used for all experiments, because the monkeys had great difficulty making accurate saccades to multiple auditory targets, even after months of training. For details on training procedures, see Chapter 2.

In both tasks, trials began with the appearance of a fixation light, usually directly in front of the monkey at (0°, 0°). (For two units recorded in areas that were clearly responsive to downward saccades and to stimuli in the lower hemifield, the fixation

Memory-saccade Task



Fixation Task

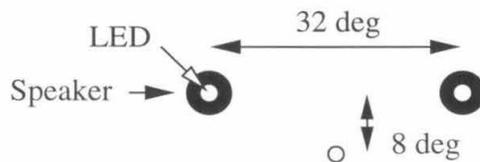
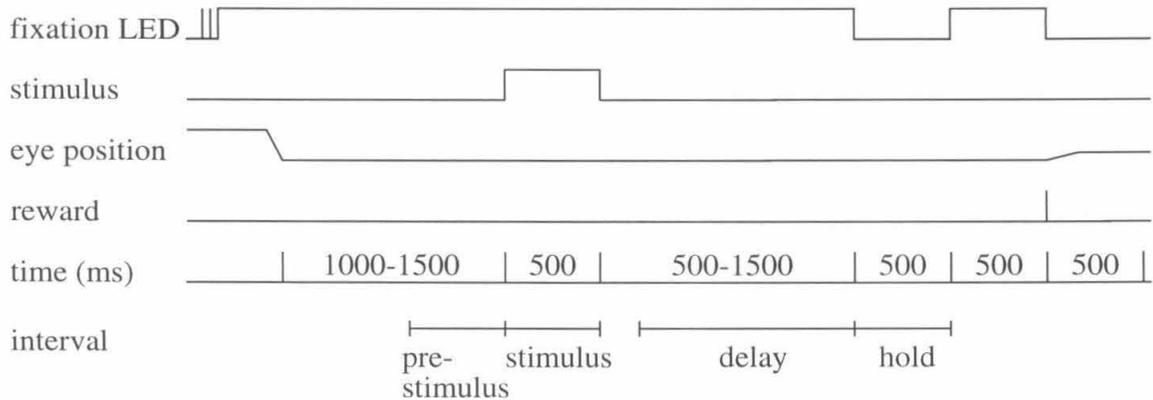


Figure 3.1: Diagrams of the memory-saccade task, fixation task, and stimulus array. In the pre-stimulus, stimulus, and delay intervals, the behavioral requirements of the two tasks were identical. In the saccade/hold interval, the behavioral requirements were different; the monkey was required either to make a saccade (in the memory-saccade task), or to hold fixation in darkness (in the fixation task). A task cue was provided with the fixation light, which remained steady after onset in memory-saccade trials, but flashed twice at the start of fixation trials. See text for details.

light was positioned at (0° , $+16^\circ$), above the two stimulus locations.) The fixation light remained steady after onset in the memory-saccade task, but flashed on and off for 200 ms (and then stayed on) at the beginning of the fixation task. This flash cue was provided to indicate to the animal which type of task he was expected to perform on a given trial. The monkey was required to fixate the central light within 1 s of its appearance, and to hold his eye position within a circular window of radius $2\text{--}3^\circ$ centered on that light. After a 1000–1500 ms interval, an auditory or visual stimulus appeared for 500 ms at one of two possible stimulus locations: left (-16° , $+8^\circ$) or right ($+16^\circ$, $+8^\circ$). The fixation light remained illuminated through this 500 ms stimulus presentation period and through a variable delay period after stimulus offset. For the majority of the experiments, the delay period was 1000–1500 ms; in the earliest experiments, a 500–1000 ms or 800–1300 ms delay period was used. The monkey was required to maintain fixation through the stimulus and delay periods in both the memory-saccade and the fixation tasks. Except for the flashing LED at the start of fixation trials, all differences between the two tasks occurred after the fixation light was extinguished.

In the memory-saccade task, the monkey was required to make a saccade within 500 ms after fixation light offset, to bring his eye position into an $8\text{--}16^\circ$ -radius window centered $0\text{--}6^\circ$ above the location at which the auditory or visual stimulus had earlier appeared. Eye position window parameters were adjusted within this range for each monkey to accommodate individual variability in memory-saccade trajectories. As previous studies have shown (Gnadt et al. 1991; White et al. 1994), visual memory saccades display a characteristic upshift, and are far more variable in endpoint than visually guided saccades. Auditory memory saccades recorded in the present study showed comparable upshift and endpoint variability, but were slightly larger in total amplitude (and, for monkey B, slower in both latency and peak speed) than visual memory saccades made under identical behavioral conditions.

After completing a memory saccade, the monkey was required to hold his eyes within the eye position window for 500 ms. Then, an LED was illuminated at the true target location. To complete the memory-saccade trial and receive a reward, the

monkey was required to make a corrective saccade to this visual stimulus within 100–250 ms, and to hold his eye position for 500 ms within a 4° -radius window centered on the visual stimulus.

In the fixation task, the monkey was required to continue fixating straight ahead in total darkness after fixation light offset. The animal had to keep his eye position steady for 500 ms within a 4° -radius window centered on the fixation point. Then, the fixation light was re-illuminated. The monkey's eye position was required to be within a $2\text{--}3^\circ$ -radius window around the fixation light within 50 ms of its re-appearance; after holding his eye position steady on the re-illuminated fixation light for 500 ms, the animal received a reward. The time course of the fixation task was therefore very similar to the time course of the memory-saccade task, except that the animal was required to hold fixation, not to make a saccade, when the fixation light was extinguished. Eye position was recorded for at least 500 ms after the reward, so that very late saccadic eye movements could be monitored.

All behavioral requirements, including eye position window parameters, were identical for auditory and visual trials of the same task. Moreover, auditory and visual stimulus presentations at the left and right stimulus locations were always interleaved (and presented in a balanced pseudorandom order, so that each of the four trial conditions appeared at least once in every set of ten successful trials for each task). The monkey was rewarded with a drop of water or juice for fulfilling all of the behavioral conditions in a given trial. The success rate for memory-saccade trials was usually 80–90%. The success rate for fixation trials was usually above 90%.

3.2.4 Recording procedures and strategy

Details of the recording procedures are described in Chapter 2. Briefly, single-unit extracellular recording was performed using tungsten microelectrodes, and all penetrations were approximately normal to the gyral surface. To help ensure that recordings came from area LIP (within the intraparietal sulcus) rather than area 7a (on the gyral surface), the electrode was advanced to 2500–3000 μm below the dura be-

fore any data were collected. Eye position was monitored using the scleral search coil technique (Judge et al. 1980) and recorded at 1000 samples/sec. At the start of each recording session, the animal was required to fixate visual stimuli at each of the stimulus locations used in the experiment, and eye position recording equipment was calibrated.

Monkeys performed the auditory and visual memory-saccade tasks described above while the recording electrode was advanced in search of neurons. Once a neuron had been isolated, data were collected during a complete block (approximately 10 trials per condition) of interleaved auditory and visual memory-saccade trials. In each trial, an auditory or visual stimulus appeared at one of the two possible stimulus locations, (-16° , $+8^\circ$) or ($+16^\circ$, $+8^\circ$); locations of auditory and visual stimuli were not optimized for the cell's receptive field. If the neuron seemed (by visual inspection of responses) to show modulation of its response in *any* period of either the auditory or the visual memory-saccade task, data collection continued with a block of interleaved auditory and visual fixation trials, during which stimuli were presented at the same two locations. Memory-saccade trial blocks were alternated with fixation trial blocks for as long as the isolation could be maintained. Typically, one or two blocks were recorded for each task, with 10 trials per condition in each block.

3.2.5 Database

The database consists of 160 unit recordings (99 neurons from monkey B, left hemisphere; 61 neurons from monkey Y, right hemisphere) for which data were collected during at least one block of memory-saccade trials and one block of fixation trials. As explained previously, the animals performed blocks of memory-saccade trials and blocks of fixation trials in alternation during each recording, for as long as the neural isolation seemed stable. Most of the recordings (134 neurons) include equal numbers of memory-saccade and fixation blocks (79 neurons, one block of each task; 54 neurons, two blocks of each task; 1 neuron, three blocks of each task). The remaining few recordings (26 neurons) ended after the second memory-saccade trial block, and

therefore include two memory-saccade blocks and one fixation block. Trials involving auditory and visual (and left and right) stimulus presentations were interleaved within each task block.

3.2.6 Data analysis

Unless noted otherwise, analyses are conducted on data pooled across monkeys; all significant results for pooled data are significant in data for the first monkey (monkey B) alone, and either significant or evident as a consistent trend in data for the second monkey (monkey Y, from whom fewer cells were recorded). Since pooled data combine recordings made from different hemispheres in the two monkeys, stimulus locations are identified throughout the text as contralateral or ipsilateral, relative to the hemisphere in which recordings were made. All analyses involve comparison of mean firing rates between contralateral trials (trials involving contralateral stimulus presentations) and ipsilateral trials (trials involving ipsilateral stimulus presentations). Only differences between contralateral and ipsilateral trials are analyzed, because changes in firing rate which are equivalent for contralateral and ipsilateral trials cannot be distinguished from general arousal effects. However, the trends discussed in this chapter persist even when such non-specific responses are also considered.

Neural responses are analyzed in four different intervals: the *pre-stimulus* period (the 500 ms interval before auditory or visual stimulus onset), the *stimulus* period (the 500 ms interval from stimulus onset to stimulus offset), the *delay* period (the 300–1300 ms interval extending from 200 ms after stimulus offset to fixation offset), and the *saccade/hold* period (the 500–800 ms interval from fixation offset to onset of the corrective visual cue). Note that the animal’s behavior during the pre-stimulus, stimulus, and delay periods was identical in the memory-saccade and fixation tasks. During the saccade/hold period, the animal either made a saccade (in the memory-saccade task) or held his eye position steady without a fixation point (in the fixation task). All analyses are based on correctly completed trials from neural recordings which included at least one block of memory-saccade trials and at least one block of

fixation trials.

Analyses of *response differentials* in a given period involve, for each neuron in the population, calculation of the difference between the mean firing rate in that period during contralateral trials and the mean firing rate during ipsilateral trials. The response differential is therefore the component of the neuron’s response which varies with stimulus location, a measure of spatial tuning. An individual neuron has a *significant spatially tuned response* (or a *significant response differential*) in a given period if there is a significant difference in mean firing rate between contralateral and ipsilateral trials during that period (Mann-Whitney test, significance level 0.05).

Throughout the text, firing rates and response differentials are expressed in spikes per second (Hz), and nonparametric analysis methods are used wherever possible. All statistical tests are two-tailed, and the critical significance level is 0.05 (n.s. means “not significant at the 0.05 significance level”). Applications of bootstrap methods involve 1000 iterations; in each iteration, a new bootstrap data set is constructed from the original data set by sampling with replacement.

3.2.7 Histology

Electrolytic lesions were placed at two penetration sites in monkey B at the end of these experiments. Histological reconstruction of these lesion sites, described in Chapter 2, indicated that the electrode penetrations were made in the lateral bank of the intraparietal sulcus. Monkey Y is still a subject in ongoing experiments.

3.3 Results

3.3.1 Behavioral modulation: stimulus period

Many neurons recorded in area LIP responded more strongly to auditory stimuli during the memory-saccade task than during the fixation task. Figure 3.2 displays the activity of an LIP neuron during presentations of auditory stimuli at the contralateral and ipsilateral stimulus locations, in the memory-saccade task and in the fixation task.

Like several other neurons in the database, this neuron has a spatially tuned auditory response; the contralateral auditory stimulus evokes significantly stronger firing than the ipsilateral auditory stimulus in both tasks (Mann-Whitney test on mean firing rates in the stimulus period: memory-saccade task $p < 0.001$, fixation task $p < 0.05$). Moreover, like other neurons in the database, this cell is more strongly activated by auditory stimuli in the memory-saccade task than in the fixation task.

In contrast, most visually responsive neurons recorded in area LIP responded similarly in the memory-saccade and fixation tasks. Figure 3.3 shows the activity of an LIP neuron during presentations of visual stimuli. This neuron has a spatially tuned visual response in both tasks; the mean firing rate in the stimulus period is significantly higher for contralateral trials than for ipsilateral trials (Mann-Whitney test, $p < 0.001$ for both tasks). However, unlike the spatially tuned auditory response of the neuron in Figure 3.2, the spatially tuned visual response of this cell appears almost equally strong in the memory-saccade and fixation tasks.

Behavioral modulation of auditory and visual responses across the population is illustrated in Figure 3.4. The four plots in this figure show response differentials (differences in mean firing rate between contralateral and ipsilateral trials) for the fixation task plotted against response differentials for the memory-saccade task, for the stimulus and pre-stimulus periods of both auditory and visual trials. Of the 160 cells in each stimulus-period plot (Figure 3.4 *A* and *B*), response differentials are significant for 35/160 (22%) in the auditory memory-saccade task, 89/160 (56%) in the visual memory-saccade task, 19/160 (12%) in the auditory fixation task, and 73/160 (46%) in the visual fixation task. Closed circles (\bullet) represent cells with a significant response differential in at least one of the two tasks; open circles (\circ) represent cells for which neither response differential is significant. All 160 neurons in the database are included in this figure, so that an unbiased estimate of behavioral modulation across the population can be obtained. Since many of the neurons have no spatially tuned response (because stimulus locations were not optimized for each cell), a large cluster appears near the origin in all four plots.

Behavioral modulation is assessed in two ways for the data in each plot. First, the

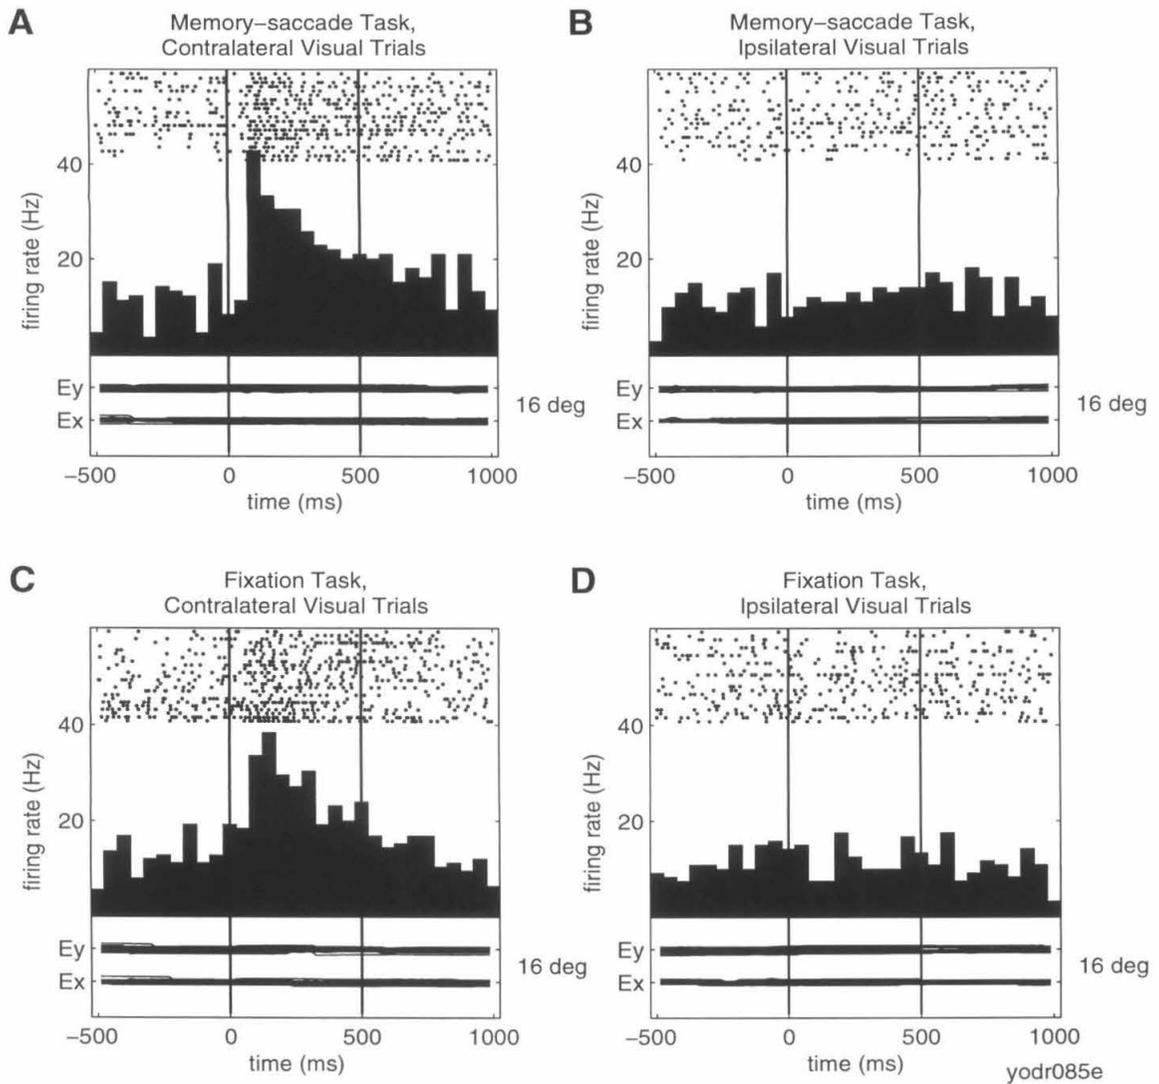


Figure 3.3: Activity of an LIP neuron during presentations of visual stimuli in the memory-saccade and fixation tasks. Conventions are the same as in Figure 3.2. The visual response of this neuron is spatially tuned (Mann-Whitney test, $p < 0.001$ for both tasks) and very similar in the two tasks (response differentials: memory-saccade task 13.4 Hz, fixation task 13.9 Hz).

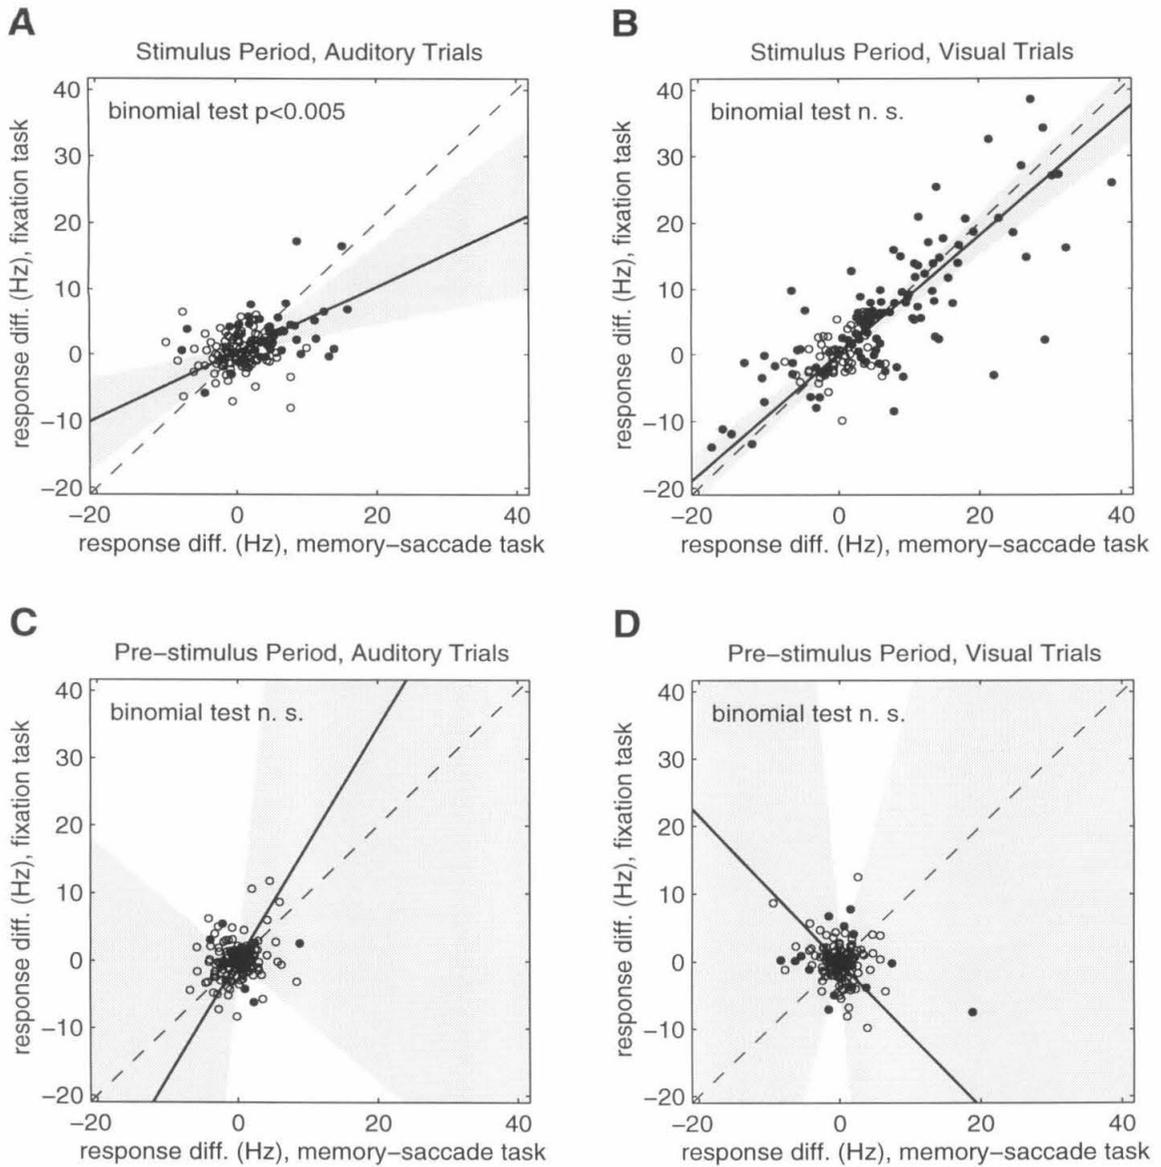


Figure 3.4: Effects of behavioral task on spatial tuning in the stimulus and pre-stimulus periods, for all 160 cells in the database. See text for explanation of plots. *A* and *B*: Response differentials in the stimulus period are significantly modulated by behavioral task for auditory trials (*A*), but not for visual trials (*B*). *C* and *D*: Response differentials in the pre-stimulus (background) period are not affected by behavioral task for either auditory (*C*) or visual (*D*) trials. Slopes of the best-fit lines, and 95% confidence intervals on the slopes: *A* 0.50 [0.22 0.81]; *B* 0.90 [0.77 1.04]; *C* 1.73 [-0.88 10.30]; *D* -1.08 [-12.43 3.94].

number of neurons for which the absolute value of the response differential is greater in the memory-saccade task than in the fixation task is compared to the number of neurons for which the reverse is true. (Absolute values of response differentials are used for this categorization so that excitatory and inhibitory responses are treated similarly.) Binomial test results printed on each plot indicate the significance level for rejection of the null hypothesis that equal numbers of neurons fall into the two categories; $p < 0.05$ implies significant behavioral modulation of response differentials across the population. Second, the two-dimensional least-mean-squares linear fit to the data (line minimizing sum of squared perpendicular distances to data points; i.e., direction of greatest variance in the data) is determined, and 95% confidence intervals on the slope of this line are calculated using a bootstrap technique. The solid line in each plot is the least-mean-squares linear fit; the dotted line represents unity slope; and the shaded area indicates the extent of the 95% confidence intervals. (Note that because the confidence intervals are determined through a bootstrap procedure, they are not constrained to be angularly symmetric around the best-fit line.) If the response differential in the memory-saccade task were equivalent to the response differential in the fixation task for each cell, then the slope of the linear fit would be one; this hypothesis can be rejected if the 95% confidence intervals on the slope do not include one.

These analyses reveal that responses to auditory stimuli are modulated by behavioral task. Across the population, stimulus-period response differentials for auditory trials (Figure 3.4 *A*) are significantly larger in magnitude during the memory-saccade task than during the fixation task (binomial test, $p < 0.005$; slope of best-fit line significantly less than one). In contrast, stimulus-period response differentials for visual trials (Figure 3.4 *B*) are not significantly different in the memory-saccade task and the fixation task (binomial test n.s.; slope of best-fit line not significantly different from one). Behavioral modulation of visual responses is therefore weak or nonexistent. (Some evidence for weak behavioral modulation of visual responses does exist in the data; while behavioral modulation of visual responses is not significant for either monkey individually according to the binomial test, the slope of the best-fit line is

significantly below one for monkey Y.) For comparison, response differentials in the pre-stimulus period are presented in Figure 3.4 *C* and *D*. The pre-stimulus period response differentials are not significantly modulated by task during either auditory or visual trials (binomial tests n.s.; slopes not significantly different from one).

The data in Figure 3.4 *A* cover a smaller range than the data in Figure 3.4 *B*, indicating that response differentials in the stimulus period are generally weaker during auditory trials than during visual trials. Could this difference in spatial tuning strength account for the apparent behavioral modulation of responses to auditory but not visual stimuli? If weakly tuned responses were modulated by task, but strongly tuned responses were not, then the analyses would indicate much more behavioral modulation for auditory trials than for visual trials. According to this explanation for the apparent behavioral modulation of auditory responses, weakly tuned visual responses should also be modulated by task. Figure 3.5, which is analogous to Figure 3.4 *B*, shows data from the 134 neurons with weak stimulus-period spatial tuning during visual trials. Neurons included in this plot have visual stimulus-period response differentials which are within the observed range of auditory stimulus-period response differentials (-10.1 to 17.2 Hz). Even for these weakly tuned neurons, no behavioral modulation of visual responses can be detected (binomial test n.s.; slope not significantly different from one in pooled data, or in each monkey's data individually). Behavioral modulation is therefore not a necessary consequence of weak spatial tuning.

These results suggest that behavioral modulation might be a distinctive characteristic of auditory responses. Another possibility, however, is that behavioral modulation might be a characteristic of auditory *cells*, rather than of auditory *responses*. In other words, the apparent behavioral modulation of auditory responses might be occurring within a small subpopulation of cells for which visual responses are also modulated by task. To address this possibility, behavioral modulation during the stimulus period was analyzed exclusively for the subpopulation of 45 *auditory cells*: cells which have significant spatially tuned responses to auditory stimuli in at least one of the two tasks. The results of this analysis (not shown) indicate that all trends

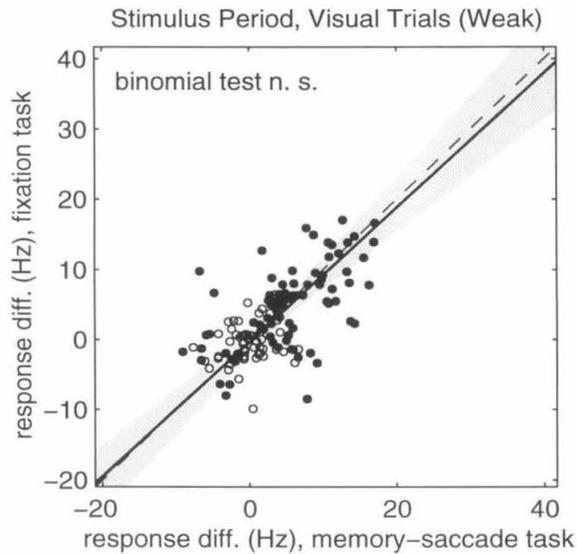


Figure 3.5: Effects of behavioral task on spatial tuning in the visual stimulus period, for cells with weak stimulus-period spatial tuning during visual trials. All neurons for which the stimulus-period response differentials for visual trials are within the observed range of stimulus-period response differentials for auditory trials (-10.1 to 17.2 Hz) are included in this plot. Conventions are the same as in Figure 3.4. No behavioral modulation of response differentials can be detected. Slope of the best-fit line, and 95% confidence intervals on the slope: 0.96 [0.79 1.16].

evident in Figure 3.4 persist when the data set is restricted to include only auditory cells. Thus, even among neurons with significant (and strongly task-dependent) auditory responses, visual responses are not significantly modulated by behavioral task. Behavioral modulation is therefore a specific characteristic of auditory responses in area LIP, rather than a general feature of both auditory and visual responses for a distinct subpopulation of LIP neurons.

3.3.2 Behavioral modulation: later periods

Many neurons recorded in area LIP responded during the delay and saccade periods of both auditory and visual memory-saccade trials, but not during the delay and hold periods of fixation trials. Figure 3.6 shows an example of stimulus-period, delay-period, and saccade-period activity recorded from a single LIP neuron during auditory and visual trials of the memory-saccade task. As in Figure 3.2, neural activity is aligned on stimulus onset. The response of this neuron is spatially tuned in the delay and saccade periods as well as in the stimulus period, for both auditory and visual memory-saccade trials (Mann-Whitney test, $p < 0.005$ for all three periods and both trial types). In the fixation task (not shown), only the response in the visual stimulus period is significantly tuned.

Across the population, spatially tuned responses tend to be stronger during the delay and saccade periods of the memory-saccade task than during the delay and hold periods of the fixation task, as illustrated in Figure 3.7. This figure is identical to Figure 3.4, except that response differentials for the delay and saccade/hold periods are displayed instead of response differentials for the stimulus and pre-stimulus periods. Response differentials for the delay period and the saccade/hold period are significantly modulated by task in both auditory and visual trials (binomial test, $p < 0.01$ in all plots; all slopes significantly less than one). Note that behavioral modulation in the delay period (Figure 3.7 *A* and *B*) resembles behavioral modulation in the stimulus period of auditory trials (Figure 3.4 *A*). The slopes of the best-fit lines in Figure 3.7 *A* and *B* (and in Figure 3.4 *A*) are significantly less than one but greater

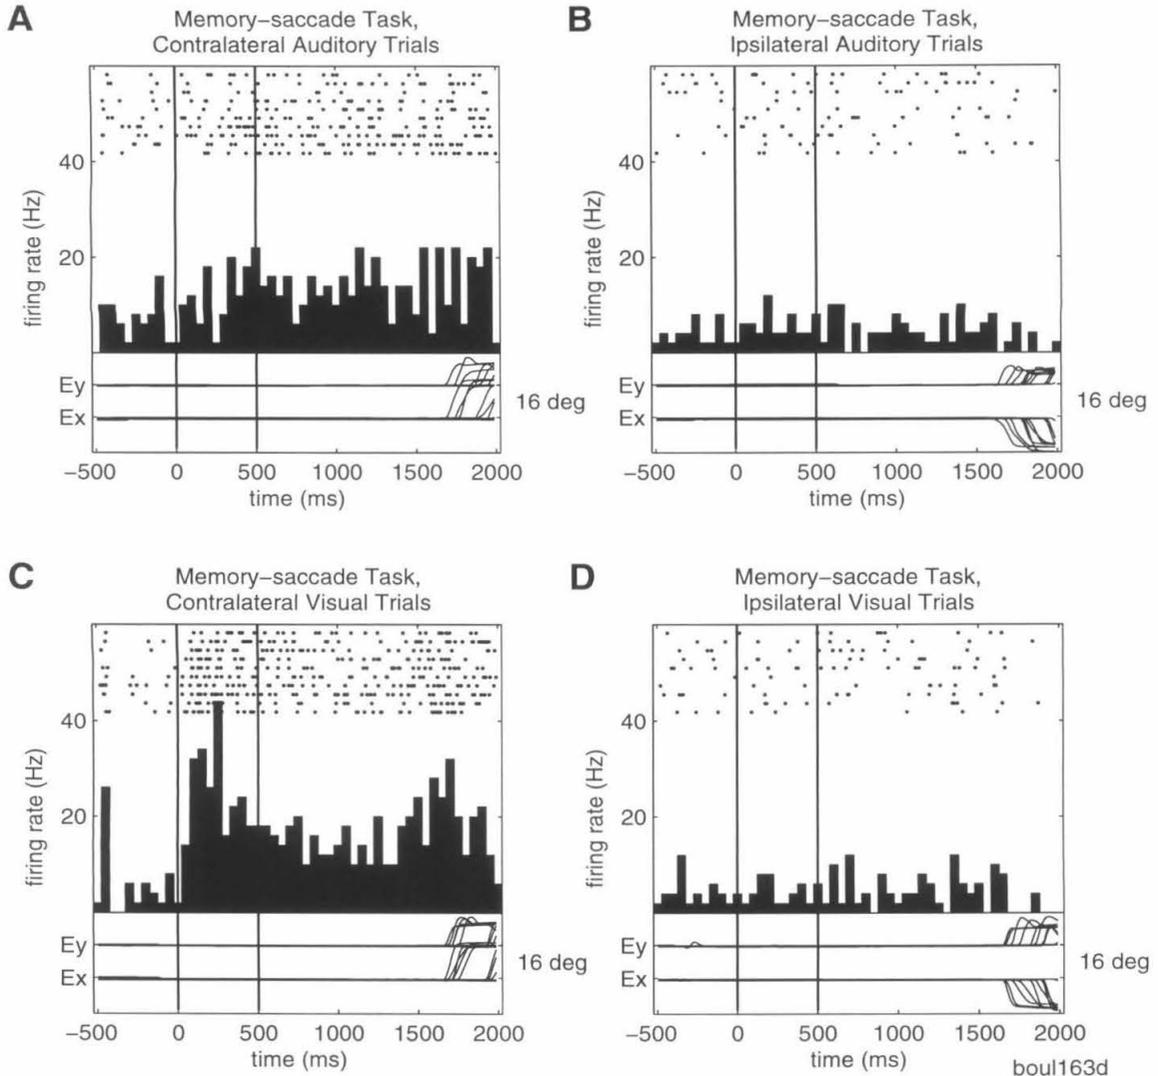


Figure 3.6: Activity of an LIP neuron during auditory and visual memory-saccade trials. Each plot shows neural activity aligned on stimulus onset; there is a variable delay of 800-1300 ms between stimulus offset and fixation offset. Large deviations in eye position traces are saccades made during the saccade period. All other conventions are as in Figure 3.2. Response differentials (with significance level) for the auditory memory-saccade task (*A* and *B*): stimulus period 6.4 Hz ($p < 0.005$); delay period 8.9 Hz ($p < 0.001$); saccade period 6.4 Hz ($p < 0.001$). Response differentials for the visual memory-saccade task (*C* and *D*): stimulus period 19.2 Hz ($p < 0.001$); delay period 10.2 Hz ($p < 0.001$); saccade period 9.9 Hz ($p < 0.001$). In the fixation task (not shown), spatial tuning is significant only during the visual stimulus period (response differentials for stimulus, delay, and hold periods of fixation task: auditory 0.6, 0.6, and -0.4 Hz; visual 18.6, 2.0, and 0.9 Hz).

than zero, while the slopes in Figure 3.7 *C* and *D* are not significantly greater than zero.

As the slopes in Figure 3.7 *A* and *B* suggest, response differentials in the memory-saccade task and the fixation task are significantly correlated in the delay period for both auditory trials (Spearman rank correlation coefficient $rs = 0.23$, $p < 0.005$) and visual trials ($rs = 0.40$, $p < 0.001$). Note that in the delay period, the only difference between the two tasks is the presumed behavioral state of the animal. In the memory-saccade task, the monkey is assumed to be remembering the location of the stimulus and planning an eye movement, while in the fixation task, the monkey is assumed to be concentrating on fixating. If these assumptions were incorrect — if, for instance, the monkey were planning to make a memory saccade after the reward in the fixation task — then response differentials in the delay period of the fixation task might be correlated with response differentials in the delay period of the memory-saccade task. In other words, one possible explanation for the correlation between memory-saccade and fixation response differentials during the delay period is that the monkeys interpreted the fixation task as an unusually complicated, very-long-delay version of the memory-saccade task.

One piece of evidence against this hypothesis is that correlation between the two tasks is much weaker in the saccade/hold period ($rs = -0.08$, n.s. for auditory trials; $rs = 0.17$, $p < 0.05$ for visual trials). If the monkeys were making saccades after the reward in the fixation task, correlation between the two tasks should have persisted in the saccade/hold period, since neural activity associated with saccade preparation should have appeared in both the saccade period of the memory-saccade task and the hold period of the fixation task. The relatively weak response correlation in the saccade/hold period might therefore be interpreted as an indication that the monkeys were *not* planning memory saccades after the reward in the fixation task. However, since the behavioral requirements of the two tasks are different in the saccade/hold period, it is conceivable that response correlation might decrease in that period regardless of the monkey's behavior after the reward.

The possibility still remains, then, that delay-period correlations might arise be-

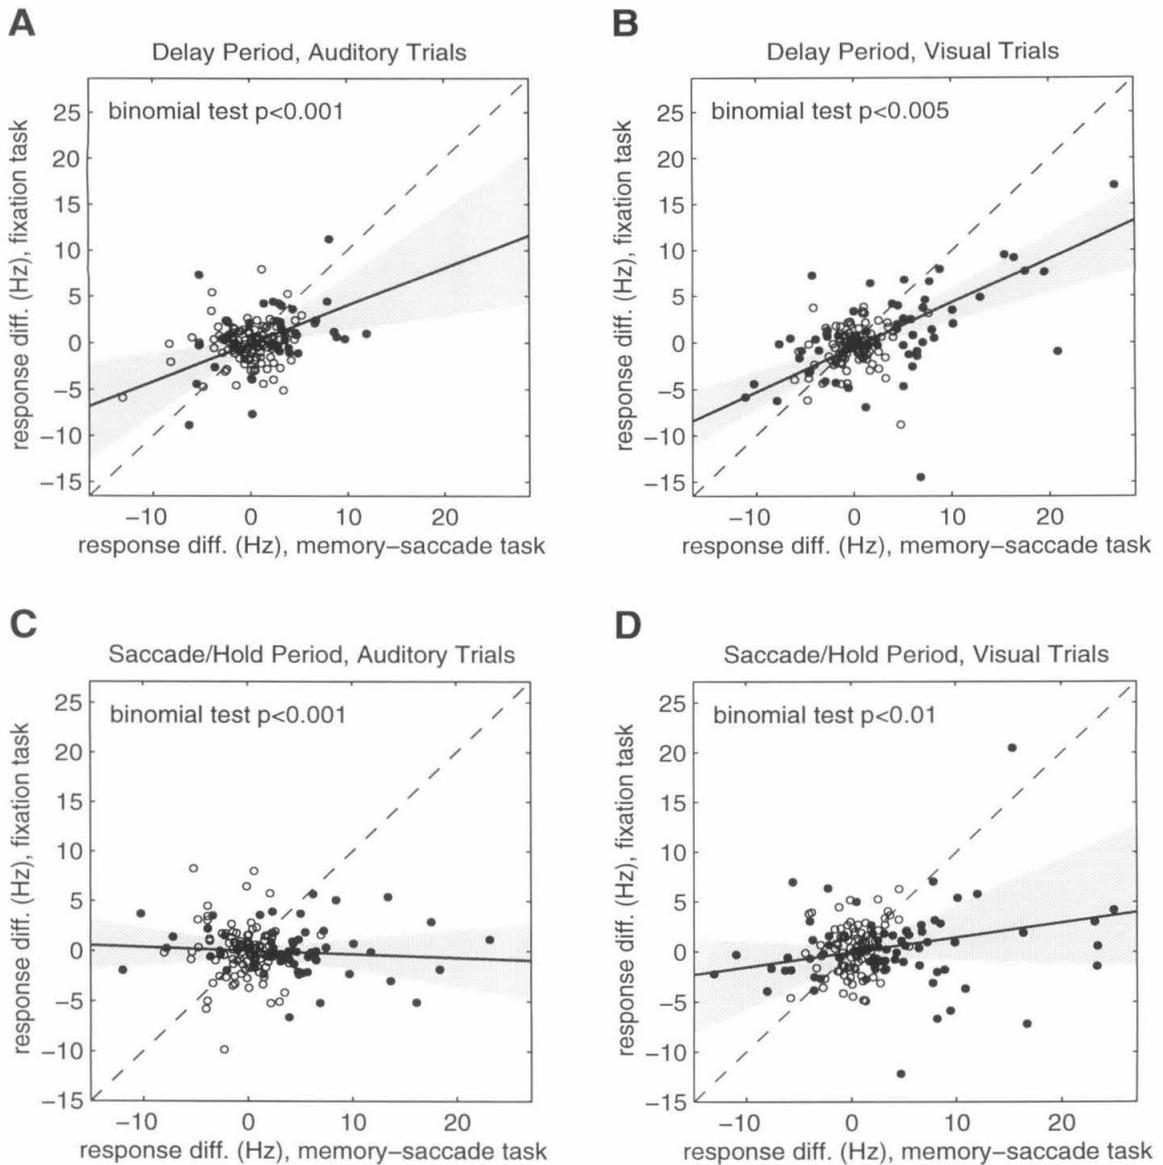


Figure 3.7: Effects of behavioral task on spatial tuning in the delay and saccade/hold periods. Conventions are the same as in Figure 3.4. *A* and *B*: Response differentials in the delay period are significantly modulated by behavioral task during both auditory (*A*) and visual (*B*) trials. *C* and *D*: Behavioral modulation is also significant in the saccade/hold period, for both auditory (*C*) and visual (*D*) trials. Slopes of the best-fit lines, and 95% confidence intervals on the slopes: *A* 0.41 [0.15 0.73]; *B* 0.48 [0.31 0.61]; *C* -0.04 [-0.19 0.09]; *D* 0.15 [-0.04 0.48].

cause the monkeys made memory saccades after the reward in the fixation task. To address this possibility directly, eye position was recorded after the reward in every fixation trial, and saccadic eye movements were identified using eye velocity criteria (optimized by visual inspection of eye traces). In the majority of fixation trials, the monkey did indeed make a single saccade within 500 ms after the reward. However, these post-reward eye movements did not appear to be directed toward the stimulus locations. Post-reward saccades, when they occurred, were similar for contralateral and ipsilateral trials, and seemed to be highly stereotyped movements toward a default eye position slightly off the fixation point.

To quantify these observations, eye positions at the end of the first post-reward saccade (or at the end of the post-reward recording period, for trials in which no saccade could be detected) were analyzed separately for every neural recording in the database. Recordings for which horizontal eye position distributions after the reward differed significantly between contralateral and ipsilateral fixation trials (Kolmogorov-Smirnov test, significance level 0.05) were judged to be contaminated by possible goal-directed movements. By this test, possible goal-directed eye movements occurred after auditory fixation trials in 6 out of 160 recordings, and after visual fixation trials in 31 out of 160 recordings. When these potentially problematic recordings are excluded from further consideration, memory-saccade and fixation response differentials are still significantly correlated in the delay period ($rs = 0.22$, $p < 0.01$ for auditory trials in reduced dataset; $rs = 0.41$, $p < 0.001$ for visual trials in reduced dataset). Therefore, the observed correlation between delay activity in the memory-saccade task and delay activity in the fixation task cannot be attributed to overt post-reward eye movements in the fixation task. It is possible, however, that goal-directed eye movements might be planned in the delay period of the fixation task but then cancelled in the hold period.

3.3.3 Auditory-visual correlation: stimulus period

Like the cell shown in Figure 3.7, many neurons recorded in area LIP responded to both auditory and visual stimuli in at least one of the two tasks. The association between auditory and visual responses across the population is illustrated for each task in Figure 3.8. Each panel in this figure shows the response differential during auditory trials plotted against the response differential during visual trials, for all 160 neurons in the database. Closed circles (\bullet) represent cells with a significant response differential in either auditory or visual trials; open circles (\circ) represent cells for which neither response differential is significant. The Spearman rank correlation coefficient r_s is indicated on each plot, along with the significance level for rejection of the null hypothesis (no correlation). As in previous scatter plots, the solid line in each plot is the two-dimensional least-mean-squares linear fit to the data, and ninety-five percent confidence intervals on the slope are indicated in gray; in these plots, however, the dotted line represents zero slope.

Response differentials in the stimulus period of auditory trials are significantly correlated with response differentials in the stimulus period of visual trials (Figure 3.8 *A* and *B*) for both the memory-saccade task ($r_s = 0.38$, $p < 0.001$) and the fixation task ($r_s = 0.25$, $p < 0.005$). The correlation coefficients for both tasks are not only significantly different from zero but also positive, indicating that the direction of spatial tuning tends to be similar for auditory and visual responses recorded from the same neuron. The low slopes of the best-fit lines in Figure 3.8 *A* and *B* confirm earlier observations that responses to auditory stimuli are generally weaker than responses to visual stimuli. For comparison, response differentials for the pre-stimulus period are shown in Figure 3.8 *C* and *D*; no correlation between auditory and visual trials is evident in the pre-stimulus period for either task ($r_s = -0.03$, n.s. for both tasks).

Further evidence that auditory responses tend to be associated with visual responses emerges from the anatomical distribution of neurons with auditory or visual responses. Figure 3.9 shows the distribution across electrode penetration sites of neurons with significant spatially tuned auditory or visual responses in the stimulus

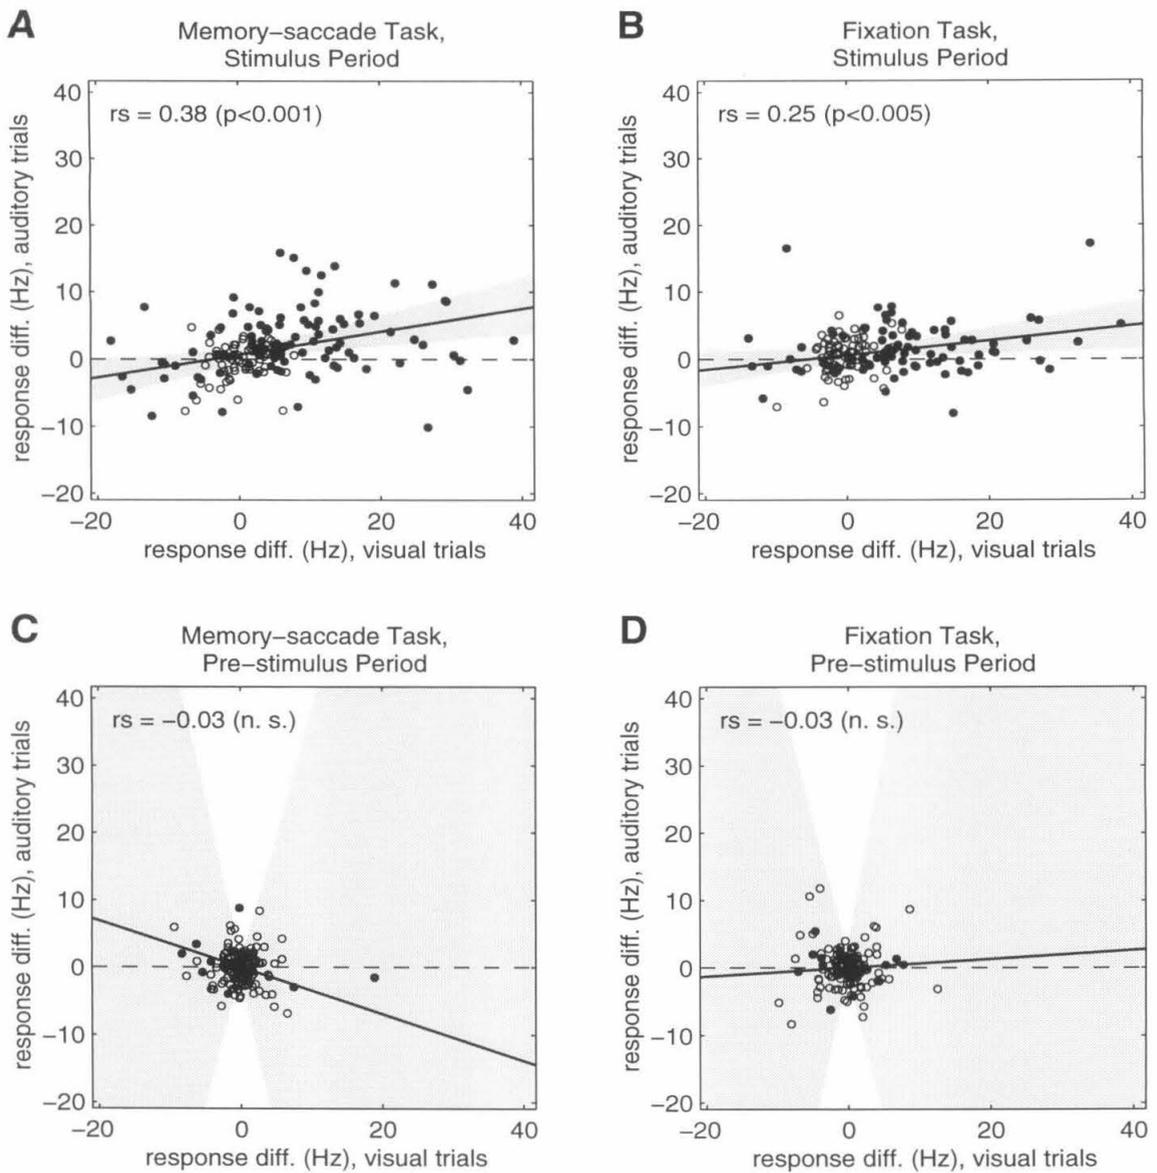


Figure 3.8: Effects of trial modality on spatial tuning in the stimulus and pre-stimulus periods, for all 160 neurons in the database. See text for explanation of plots. *A* and *B*: During the stimulus period, auditory and visual response differentials are significantly correlated in both the memory-saccade task (*A*) and the fixation task (*B*). *C* and *D*: No significant correlation between auditory and visual trials can be detected in the pre-stimulus period in either task. Slopes of the best-fit lines, and 95% confidence intervals on the slopes: *A* 0.17 [0.07 0.28]; *B* 0.11 [0.02 0.20]; *C* -0.35 [-4.88 3.89]; *D* 0.06 [-4.14 5.65].

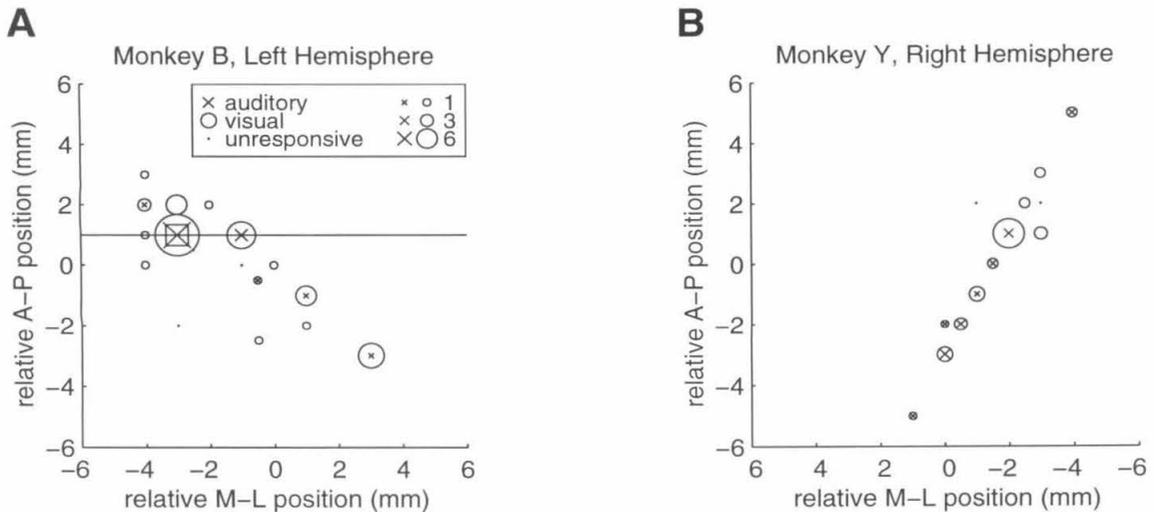


Figure 3.9: Distribution across electrode penetration sites of neurons with significant spatially tuned auditory or visual responses in the stimulus period of the memory-saccade task. Conventions are as in Figure 2.9. *A*: Electrode penetration sites for monkey B, whose recording chamber was mounted over the left hemisphere. The square shows the site of one of the electrolytic lesions made in this animal, and the line indicates the approximate angle of the histological section shown in Figure 2.2. *B*: Electrode penetration sites for monkey Y, whose recording chamber was mounted over the right hemisphere. For both monkeys, cells with auditory or visual responses are intermingled across penetration sites.

period of the memory-saccade task. (A similar figure in Chapter 2 shows the distribution of neurons with significant spatially tuned auditory or visual responses in the stimulus period of the fixation task.) In both monkeys, all penetration sites which produced cells with spatially tuned auditory responses also produced cells with spatially tuned visual responses. Moreover, neurons with auditory responses and neurons with visual responses are distributed across all the penetration sites, with no evident clustering. This overlap of auditory and visual data across penetration sites suggests that neurons with spatially tuned responses to auditory stimuli are well integrated with visually responsive neurons across area LIP.

3.3.4 Auditory-visual correlation: later periods

Correlation between auditory and visual trials occurs in the delay and saccade periods of the memory-saccade task, as well as in the stimulus period. Across the population of recorded cells, response differentials for auditory and visual trials are significantly correlated in the delay ($rs = 0.57, p < 0.001$) and the saccade ($rs = 0.66, p < 0.001$) periods of the memory-saccade task (Figure 3.10 *A* and *C*). Like the stimulus-period correlation coefficients, these delay- and saccade-period correlation coefficients are not only significantly different from zero but also positive, indicating consistent spatial tuning for delay/saccade activity recorded from the same neuron during auditory and visual memory-saccade trials. No significant correlation between auditory and visual trials is evident in response differentials for either the delay period or the hold period of the fixation task (Figure 3.10 *B* and *D*).

3.3.5 Correlation between periods

Studies of visual responses in area LIP have noted that many visually responsive neurons are active in the delay or saccade periods of a memory-saccade task (Barash et al. 1991a). Are cells with auditory responses even more likely to exhibit delay or saccade activity than cells with visual responses? Since auditory responses tend to co-occur with visual responses, this question is best addressed through comparison of two populations of neurons selected to be distinct: those with significantly tuned auditory (and possibly visual) responses in the stimulus period, and those with significantly tuned visual but *not* auditory responses. In the memory-saccade task, 66% (23/35) of neurons with spatially tuned auditory responses in the stimulus period also have delay-period responses, while 39% (25/64) of neurons with exclusively visual stimulus-period responses are active during the delay period. Thus neurons with auditory stimulus-period responses are significantly more likely than neurons with exclusively visual stimulus-period responses to exhibit delay activity (Fisher-Irwin test, $p < 0.05$). Delay-period responses were pooled across auditory and visual trials to obtain the above results; however, significant associations between auditory responses and delay

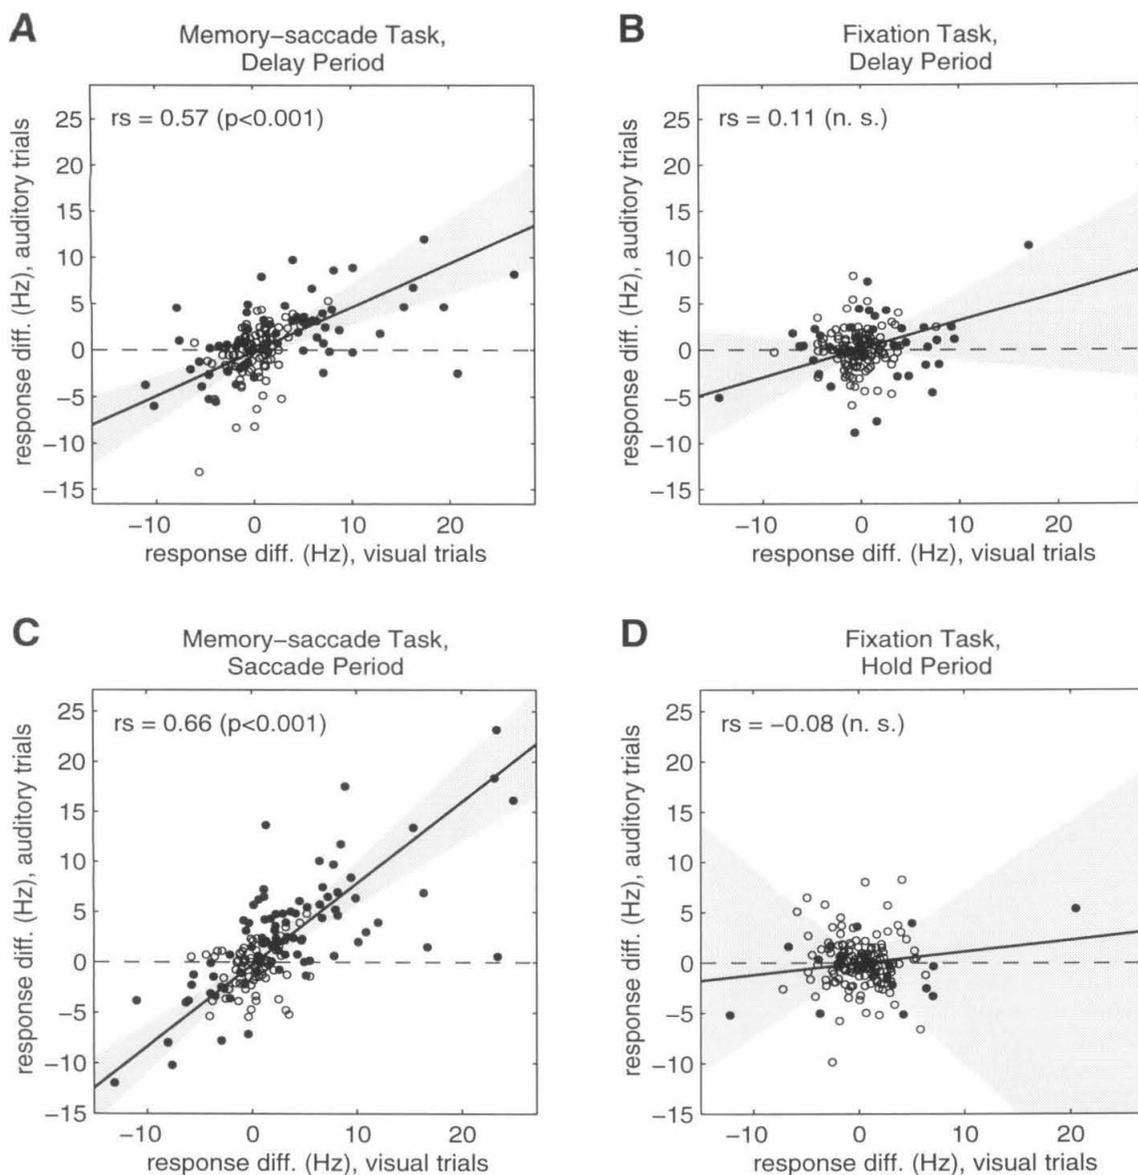


Figure 3.10: Effects of trial modality on spatial tuning in the delay and saccade/hold periods. Conventions are the same as in Figure 3.8. *A* and *C*: Response differentials during auditory and visual trials are significantly correlated in the delay (*A*) and saccade (*C*) periods of the memory-saccade task. *B* and *D*: No significant correlation between auditory and visual trials can be detected in the delay (*B*) and hold (*D*) periods of the fixation task. Slopes of the best-fit lines, and 95% confidence intervals on the slopes: *A* 0.47 [0.32 0.70]; *B* 0.30 [-0.11 0.59]; *C* 0.81 [0.62 1.01]; *D* 0.12 [-0.96 0.71].

activity are also found when delay-period responses are considered separately for auditory and visual trials.

Results for the saccade period are similar. Over 77% (27/35) of neurons with auditory stimulus-period responses in the memory-saccade task respond during the saccade period, while only 52% (33/64) of exclusively visual cells respond during the saccade period. Neurons with auditory responses in the stimulus period are therefore significantly more likely to show saccade activity than neurons with exclusively visual stimulus-period responses (Fisher-Irwin test $p < 0.05$). Again, this trend is evident not only when saccade-period responses are pooled across auditory and visual trials, but also when auditory and visual trials are considered separately.

These results indicate that auditory responses in the stimulus period of the memory-saccade task are more closely linked to delay and saccade activity than are exclusively visual responses. Could auditory responses be used to identify a subpopulation of visually responsive neurons in area LIP which are likely to be active in later phases of the memory-saccade task? To find out, two populations of visually responsive neurons can be compared (Figure 3.11): *bimodal* cells, defined to be neurons with spatially tuned stimulus-period responses during both visual and auditory memory-saccade trials; and *unimodal* (exclusively visual) cells, defined to be neurons with spatially tuned stimulus-period responses during visual but not auditory memory-saccade trials.

All four panels of Figure 3.11 show data taken from visual trials of the memory-saccade task; Figure 3.11 *A* and *C* display data from bimodal cells in the database, while Figure 3.11 *B* and *D* display data from unimodal visual cells. The division of visually responsive neurons between the left and right halves of the figure is therefore determined entirely by the presence or absence of auditory responses. Other conventions are as in Figure 3.8. (The gray areas representing the slope confidence intervals have a distinct bow-tie shape in this figure, because the bootstrapped fit lines from which the confidence intervals were determined varied in intercept; intercept variation is also present, but not so noticeable, in previous figures.) As shown in the figure, the correlation between stimulus-period response differentials and delay-period response differentials during visual trials is much stronger for neurons with

both auditory and visual stimulus-period responses than for neurons with exclusively visual stimulus-period responses (bimodal cells: $rs = 0.70$, $p < 0.001$; unimodal visual cells: $rs = 0.20$, n.s.). The difference between the two correlation coefficients is significant (Fisher z -transformation test, $p < 0.01$), and the slope of the best-fit line in Figure 3.11 *A* is significantly greater than the slope of the best-fit line in Figure 3.11 *B*. The distinction between bimodal and unimodal visual cells is weaker in the saccade period (Figure 3.11 *C* and *D*); although the correlation coefficient is slightly larger and the slope of the best-fit line higher for bimodal cells than for unimodal visual cells, these differences are not significant.

The association between auditory responses and activity in later periods of the memory-saccade task suggests that auditory responses themselves might be saccade-related. Analysis of error trials — memory-saccade trials in which the monkeys made saccades to the incorrect location — could, in principle, be used to determine whether auditory responses are in fact more dependent on the upcoming saccade trajectory than on the auditory stimulus location. Unfortunately, the statistical power of error trial analysis was very low for this data set, because there were few error trials. Comparison of stimulus-period response differentials for error trials with stimulus-period response differentials for correct trials revealed neither significant anti-correlation nor significant correlation, and was therefore inconclusive. Analysis of possible relationships between auditory responses and saccade parameters in correct trials was also inconclusive.

3.3.6 Control for response measure

Raw response differentials reflect the magnitude of spatial tuning, a quantity which is only indirectly related to the significance of spatial tuning. Analyses of response-differential distributions (Figures 3.4, 3.5, 3.7, 3.8, 3.10, and 3.11) might therefore overemphasize data from high-firing but poorly tuned cells. To control for possible artifacts associated with the use of raw response differentials, all analyses of response-differential distributions were repeated using three different normalized response mea-

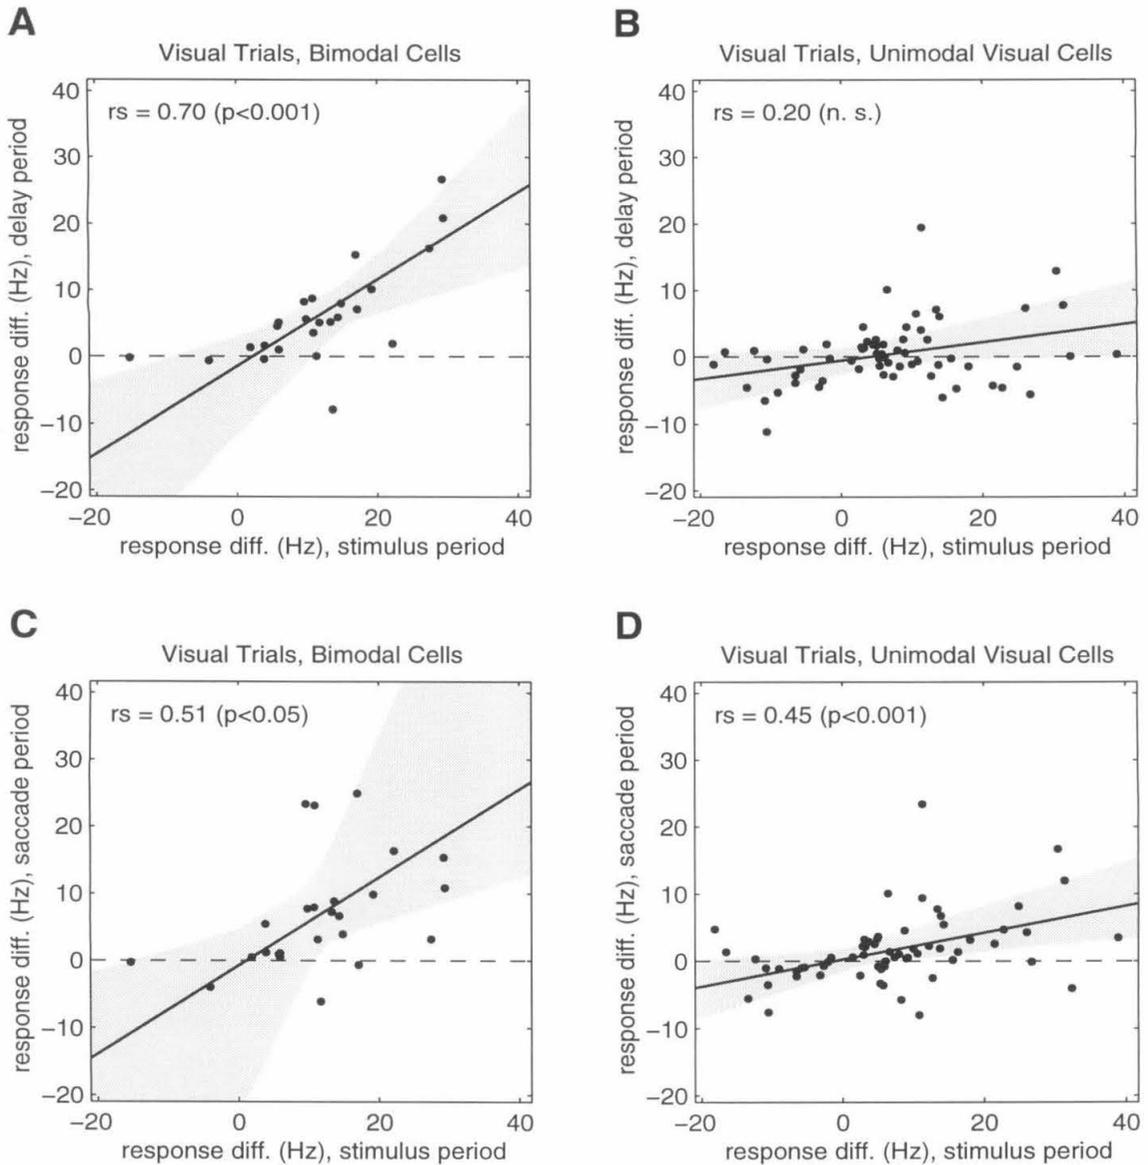


Figure 3.11: Relationship between spatial tuning in the delay or saccade periods and spatial tuning in the stimulus period, for visual memory-saccade trials only. See text for explanation of plots. *A* and *B*: Response differentials in the delay period of visual memory-saccade trials are significantly correlated with response differentials in the stimulus period for bimodal cells (*A*), but not for unimodal visual cells (*B*). *C* and *D*: Correlation between response differentials in the saccade and stimulus periods of visual memory-saccade trials is significant for both bimodal (*C*) and unimodal visual (*D*) cells. Total numbers of cells: bimodal $N = 25$ (*A* and *C*), unimodal visual $N = 64$ (*B* and *D*). Slopes of the best-fit lines, and 95% confidence intervals on the slopes: *A* 0.65 [0.32 1.14]; *B* 0.14 [0.02 0.27]; *C* 0.66 [0.29 2.38]; *D* 0.20 [0.07 0.35].

sures:

- the response differential normalized by the mean pre-stimulus-period firing rate, a measure of spatial tuning relative to background activity;
- the response differential normalized by the response sum (mean contralateral response plus mean ipsilateral response), a measure of spatial tuning relative to overall response; and
- the response differential normalized by its estimated standard error, a direct measure of the significance of spatial tuning.

Results obtained using all three normalized measures are consistent with those shown for raw response differentials.

3.3.7 Control for block order

For each neural recording in this experiment, blocks of memory-saccade and fixation trial data were always collected in the same order: first a block of memory-saccade trials, then a block of fixation trials, and so on in alternation, for as long as the isolation could be maintained. On average, then, blocks of fixation trials were collected later in each recording than blocks of memory-saccade trials. Stronger spatial tuning in the memory-saccade task than in the fixation task could, in principle, arise from systematic changes (such as a decrease in overall firing rate) over the course of each recording. One control for such effects has already been shown; response differentials in the pre-stimulus period do not appear to be modulated by task (Figure 3.4 *C* and *D*). As an additional control, response differentials for the first block of fixation trials were compared to response differentials for the second block of memory-saccade trials (for the 81 recordings with at least one block of fixation trials and two blocks of memory-saccade trials). Thus for this analysis, data were selected such that fixation blocks were collected *earlier* in each recording than memory-saccade blocks. All trends in Figures 3.4, 3.5, and 3.7 were also evident in this control analysis, confirming that observed behavioral modulation effects are not an artifact of block order.

3.4 Discussion

The main result of this chapter is that neurons in area LIP respond more strongly to auditory stimuli when monkeys are engaged in a memory-saccade task than when they are engaged in a fixation task. Additional findings are:

- Visual responses, unlike auditory responses, are not significantly modulated by behavioral task.
- Behavioral modulation of auditory responses resembles behavioral modulation of delay-period activity.
- Auditory responses are associated with visual responses in both the memory-saccade task and the fixation task.
- Auditory responses are also associated with delay or saccade activity.

Taken together, these results imply that auditory responses in area LIP are best considered supramodal (cognitive or motor) responses, rather than modality-specific sensory responses.

In combination with the results of Chapter 2, which show that auditory responses appear in the fixation task only after auditory-saccade training, these findings indicate that the last of the four possibilities raised in the introduction is correct: responses to auditory stimuli in area LIP depend both on training and on behavioral context. Therefore, the resolution to the apparent discrepancy between early studies of area LIP, which found no responses to auditory stimulation (Mountcastle et al. 1975; Hyvärinen 1982b; Koch and Fuster 1989), and later studies, which did find auditory responses in LIP (Mazzoni et al. 1996b; Stricanne et al. 1996), is that the monkeys had both learned an auditory-saccade task and been required to perform this task in the latter but not the former study. Further implications of the results, and interpretations in light of previous studies, are discussed below.

3.4.1 Behavioral modulation of auditory responses

Responses to auditory stimuli in area LIP are strongly modulated by behavior, while responses to visual stimuli do not appear to be dependent on task. Behavioral modulation of auditory responses is not a necessary consequence of weak spatial tuning, nor a general feature of all stimulus-period responses for cells which respond to auditory stimuli. Moreover, no behavioral modulation is observed in the pre-stimulus period, and behavioral modulation is not an artifact of trial block order. Behavioral modulation therefore seems to be a robust and distinctive characteristic of auditory responses in area LIP.

This study is the first to show that auditory responses in area LIP are dependent on behavioral task. However, behavioral modulation of auditory responses has previously been observed in several regions of the brain which are directly connected to area LIP. Neurons in the deep layers of the superior colliculus, for example, respond to auditory stimuli in the context of a saccade task, but habituate rapidly to auditory stimuli when no saccade is required (Jay and Sparks 1984; Jay and Sparks 1987a). Neurons in the prefrontal cortex also respond to auditory stimuli more strongly in the context of goal-directed (arm and eye) movements than in the context of an auditory detection or a passive listening task (Vaadia et al. 1986). Responses to auditory stimuli in these areas, and responses to auditory stimuli in area LIP, may best be considered cognitive or motor responses, related primarily to the significance of the stimulus as a potential target for movement.

3.4.2 No behavioral modulation of visual responses?

Across the population of neurons recorded in this study, visual responses and background (pre-stimulus) activity are not significantly modulated by behavioral task. This result seems to contradict recent reports that visual responses and background activity in area LIP are enhanced in a memory-saccade task relative to a fixation task (Colby et al. 1996). Even when re-analyzed using the analysis methods described in Colby et al. (1996), to compare maximal responses rather than response differentials

in the two tasks, the data collected in the present experiment still show no evidence for behavioral modulation of visual responses in the stimulus period (for either monkey alone or for both together), and no evidence for modulation of responses in the pre-stimulus period. The apparent discrepancies between the present study and Colby et al. (1996) are therefore not likely to be due to differences in data analysis methods.

The discrepancies between the present study and that of Colby et al. (1996) might, however, arise from differences in behavioral paradigms and recording procedures. For the present experiments, two fixed stimulus locations were used, and stimulus presentations were randomized across the two locations. The monkeys therefore did not know which of the two possible stimulus locations would be relevant on any given trial until the stimulus actually appeared. In contrast, Colby et al. (1996) optimized the stimulus location for each cell, and then used that one stimulus location for all experiments on the cell. Their monkeys therefore knew the location of the relevant stimulus even before it appeared on a given trial. Colby et al. (1996) did suggest that the background enhancement they observed in the memory-saccade task might have arisen because the monkeys were anticipating the onset of the behaviorally relevant stimulus in the receptive field. Another possibility is that enhancement of both background activity and visual responses occurred in the memory-saccade task because the monkeys were planning the impending movement (Mazzoni et al. 1996a; Bracewell et al. 1996; Shadlen and Newsome 1996; Platt and Glimcher 1997b).

3.4.3 Behavioral modulation of delay activity

Neurons in area LIP are more active in the delay and saccade periods of the memory-saccade task than in the delay and hold periods of the fixation task, for both auditory and visual trials. This result was expected. In the memory-saccade task, the monkey must remember the location of a previously presented stimulus, plan an eye movement, and execute a saccade. Delay activity is thought to reflect motor intention or spatial attention which would be engaged in the delay period of the memory-saccade task but not in the delay period of the fixation task. Similarly, saccade activity should

occur only in the saccade period of the memory-saccade task, not in the hold period of the fixation task.

A more unexpected finding is that behavioral modulation in the delay period resembles behavioral modulation in the auditory stimulus period. Like auditory responses, delay-period responses are weaker, on average, during fixation trials than during memory-saccade trials, but activity does persist in the fixation task. Indeed, response differentials in the delay period of fixation trials are significantly correlated with response differentials in the delay period of memory-saccade trials. This correlation might be considered evidence that the animals did not fully realize they were supposed to be performing a fixation task (rather than a very-long-delay version of the memory-saccade task). Certainly, delay-period activity is usually associated with movement planning or peripherally directed attention, neither of which was required in the fixation task. For three reasons, however, it seems very unlikely that the animals were misinterpreting the fixation task. First of all, the behavioral paradigm for fixation trials ensured that eye movements toward the stimulus locations within 1500–2500 ms after stimulus offset would cause the trial to be aborted. Second, the use of trial blocking and task cues (steady fixation light onset in memory-saccade trials, flashing onset in fixation trials) made the presentation of fixation trials entirely predictable. Third, the correlation does not disappear when the data set is restricted to recordings which are unlikely to be contaminated by very late, goal-directed eye movements in the fixation task.

Rather than aberrant behavioral strategies, the observed correlation in delay-period response differentials may reflect covert orienting responses or attentional effects. Auditory and visual stimuli may evoke default movement plans or sustained attentional orienting which activate area LIP during the delay period of the fixation task, even though the fixation task does not require either an eye movement or a re-direction of attention. In support of this view, previous studies have demonstrated that movement plans are represented in LIP even when the movement is never executed (Bracewell et al. 1996; Snyder et al. 1997; Snyder et al. 1998). The apparent similarity between behavioral modulation of delay-period activity and be-

havioral modulation of auditory responses therefore raises the possibility that both delay activity and auditory responses reflect default movement plans.

3.4.4 Association between auditory and visual responses

Neurons with auditory stimulus-period responses tend to have visual stimulus-period responses with similar spatial tuning, in both the memory-saccade task and the fixation task. Moreover, neurons which respond during the delay or saccade periods of auditory memory-saccade trials are likely to respond similarly during the corresponding periods of visual memory-saccade trials. No such correlation between auditory and visual trials can be detected in the delay or hold periods of the fixation task, or in the pre-stimulus period of either task. Thus, correlations between auditory and visual trials occur specifically during stimulus presentations in both tasks, and during the later phases of the memory-saccade task.

These findings are consistent with the results of previous studies of auditory and visual responses, both in area LIP and in regions of the brain which are anatomically connected to area LIP. In an earlier investigation of LIP activity during auditory and visual memory-saccade trials, Mazzoni et al. (1996b) concluded that neurons active during the stimulus, delay, or saccade periods of an auditory memory-saccade task tended to be active during the same periods of a visual memory-saccade task. The present study confirms those results, and further demonstrates that an association between auditory and visual trials also exists during the stimulus period, but not later periods, of a fixation task. Similar response correlations between auditory and visual trials, either during sensory stimulation or during later phases of a movement task, have also been noted in superior colliculus (Jay and Sparks 1984; Jay and Sparks 1987b; Wallace et al. 1996), frontal cortex (Vaadia et al. 1986), frontal and supplementary eye fields (Russo and Bruce 1994; Schall 1991a), and supplementary motor areas (Schall 1991b).

The observed correlations between auditory and visual trials during the delay and saccade periods of the memory-saccade task could be viewed as confirmation that

activity during these periods is related to target selection or movement planning. Movement cues of different sensory modalities evoke similar delay and saccade activity in LIP; therefore, this activity probably reflects supramodal processes, such as motor intention or purely spatial attention. By extension, the association between auditory and visual responses in the stimulus period implies that some component of stimulus-evoked activity in area LIP also reflects target selection or movement planning. The results therefore lend support to the idea that responses to auditory stimuli in area LIP are supramodal intentional or attentional responses, rather than modality-specific sensory responses.

3.4.5 Link between auditory and delay/saccade activity

Neurons with auditory stimulus-period responses are much more likely to display delay or saccade activity than neurons with exclusively visual stimulus-period responses. Moreover, in the visual memory-saccade task, correlation between stimulus-period and delay-period activity is higher for neurons with both auditory and visual stimulus-period responses than for neurons with exclusively visual stimulus-period responses. These findings suggest that neurons in area LIP which respond to auditory stimuli are more directly involved in eye movement planning than neurons which respond to visual stimuli alone. Given the physiological similarities between area LIP, the frontal eye fields, and the deep layers of the superior colliculus, a similar association between auditory and delay- or saccade-related activity may be evident in the frontal eye fields and the superior colliculus. Previous studies of these areas have not provided data appropriate for direct comparison with the present results.

3.4.6 Experimental considerations

The results presented in this chapter indicate that auditory responses in area LIP are dependent on behavioral task, associated with visual responses, and predictive of delay or saccade activity. It should be noted, however, that these findings (and those of the previous chapter) may be dependent on the choice of experimental conditions.

Four possible caveats seem especially worthy of consideration.

First, the auditory stimuli used in the present study were bursts of high-frequency band-limited white noise (5–10 kHz), which probably have little ethological significance for monkeys. Sounds with different spectral characteristics (e.g., macaque vocalizations) might conceivably elicit auditory responses in area LIP which are less dependent on behavioral task (or on auditory-saccade training) than the responses observed in the present study.

Second, in these experiments, auditory stimuli were presented only at locations within the visual field, at relatively small eccentricities. Since primates may use auditory spatial cues primarily for localizing targets outside of the visual field, it is possible that auditory stimuli presented at large eccentricities might evoke auditory responses in area LIP which are not associated with visual responses (nor dependent on auditory-saccade training). Moreover, if neurons in area LIP have auditory receptive fields which are more peripheral than their visual receptive fields, then the two fixed stimulus locations used in the present experiment might occasionally have been optimal for a neuron's visual receptive field, but might never have been optimal for any neuron's auditory receptive field. Apparent behavioral modulation of responses to auditory stimuli might therefore turn out to be behavioral modulation of responses to sub-optimal stimuli. This scenario seems unlikely, because weakly tuned visual responses (which presumably represent responses to sub-optimal visual stimuli) do not appear to be modulated by task (Figure 3.5); however, the possibility cannot be ruled out on the basis of the present data.

Third, the position of the pinnae was not controlled in these experiments. Therefore, the apparent link between auditory responses and eye movements might actually reflect an association between auditory responses and saccade-related pinna movements. Moreover, if the monkeys moved their pinnae differently during the stimulus periods of memory-saccade and fixation trials, apparent behavioral modulation of auditory responses might have occurred because auditory stimuli were filtered differently by the ears in the two tasks. Although these possibilities cannot be excluded, they seem very unlikely. Previous studies have shown that the incidence of auditory

responses in area LIP, and the tuning of auditory responses in superior colliculus, are not significantly altered by pinna restraint in awake monkeys (Stricanne et al. 1996; Jay and Sparks 1987a); therefore, auditory responses cannot be entirely dependent on saccade-related pinna movements. Furthermore, while pinna movements have not been studied intensively in monkeys, a recent behavioral study in cats suggests that pinna movements could not account for the observed behavioral modulation of auditory responses. Cats make auditory-evoked pinna movements, which do not appear to be dependent on behavioral task, and orienting pinna movements, which occur in conjunction with eye movements (Populin and Yin 1998). Assuming these results generalize to monkeys, pinna movements in response to auditory stimulation should have been the same for the two behavioral tasks, and pinna movements in conjunction with eye movements should not have occurred until long after the auditory stimulus period.

Finally, the monkeys used in the present study performed all the behavioral tasks with their heads immobilized. Under more natural conditions, primates orient to auditory and visual stimuli with a combined movement of the head and eyes (Whittington et al. 1981; Goldring et al. 1996). Because auditory targets can be perceived at larger eccentricities than visual targets, and can therefore evoke larger orienting movements, responses to auditory stimuli may be strongly associated with free head movement. Responses to auditory stimuli in area LIP might therefore be most robust in the context of head movements, rather than eye movements.

Although these potential caveats should not be overlooked, it seems likely that the results of the present study will generalize to other experimental conditions, because the findings are consistent with previous studies of auditory responses in areas which are anatomically connected to area LIP. In particular, behavioral modulation of auditory responses, and associations between auditory and visual responses, have been observed in both superior colliculus and frontal cortex under a range of different experimental conditions (superior colliculus: Jay and Sparks 1987a; Wallace et al. 1996; frontal cortex: Vaadia et al. 1986; Russo and Bruce 1994). The present findings are also consistent with current interpretations of LIP function, as discussed further

below.

3.4.7 Interpretations

This chapter demonstrates that responses to auditory stimuli in area LIP are dependent on behavioral task, associated with visual responses, and predictive of delay or saccade activity. These results imply that responses to auditory stimuli in area LIP are best considered supramodal responses, not modality-specific sensory responses. Several different interpretations of these findings — and of the role of area LIP in auditory-to-oculomotor transformations — are possible.

For example, auditory activity in area LIP may be related to spatial attention that is not modality specific (Colby et al. 1996; Gottlieb et al. 1998). According to this interpretation, LIP responses to auditory stimuli are stronger in the memory-saccade task than in the fixation task because the animal must attend more closely to the spatial information present in the auditory cue when a localization movement is required. The fact that auditory responses in area LIP are weaker and more dependent on behavioral task than visual responses implies, in this scenario, that auditory stimuli do not capture spatial attention as effectively as visual stimuli. Indeed, the auditory stimuli used in this experiment were probably less easy to localize (and perhaps less spatially salient) than the visual stimuli, given that the monkeys required months of training to master the auditory-saccade task but only a few days to master the visual-saccade task (see Chapter 2).

The results of the present study are also consistent with the view that activity in area LIP reflects movement intention (Mazzoni et al. 1996a; Bracewell et al. 1996; Snyder et al. 1997; Platt and Glimcher 1997b; Snyder et al. 1998). According to this interpretation, responses to auditory stimuli in area LIP are modulated by behavioral task because auditory stimuli evoke more definite movement plans in the memory-saccade task than in the fixation task; similarly, auditory responses are more task-dependent than visual responses because auditory orienting is less reflexive than visual orienting (at least for the stimuli used in this study). Residual activity

in the stimulus period of auditory fixation trials, discussed further in Chapter 2, represents a suppressed intention to make an eye movement to an auditory target made familiar by months of saccade training. Consistent with this interpretation, the link between auditory stimulus-period responses and delay or saccade activity in the memory-saccade task implies that responses to auditory stimuli in LIP are directly related to movement intention.

A third possible interpretation of the data is that responses to auditory stimuli in area LIP reflect oculomotor significance: the significance of the stimuli as potential targets for eye movements. By this argument, stimulus-period auditory activity in the fixation task reflects the learned significance of the auditory stimulus as a *possible* eye movement target. When the sound becomes an *obligate* target for an eye movement in the memory-saccade task, its significance increases further. However, in the memory-saccade task, the increase in the auditory stimulus-period response is linked to the presence of continued activity in the delay period, and other experiments have shown that delay-period activity generally reflects the monkey's intention to make eye movements (Snyder et al. 1997; Snyder et al. 1998). Thus a simpler explanation for the increase in stimulus-period activity in the auditory memory-saccade task may be that movement-planning activity is added to activity reflecting the learned significance of the auditory stimulus.

Finally, a fourth possibility is that spatial attention, movement intention, and oculomotor significance are artificial psychological distinctions for area LIP, which performs sensory-to-motor transformations for saccades. According to this view, increased activity in the stimulus period of the auditory memory-saccade task simply reflects a graded increase in the preparation for a sensory-guided eye movement.

This study was designed to resolve discrepancies between early and more recent reports regarding auditory activity in LIP, not to distinguish between the four possible interpretations of auditory activity described above. Further research will be required to determine the degree to which behavioral modulation of auditory responses supports these different interpretations. For instance, if future experiments show that auditory stimuli evoke stronger responses in LIP when a monkey plans a saccade to

an auditory target than when he plans a reach to the same target, then a significant component of auditory activity in LIP represents intention to make a saccade, independent of spatial attention. Since delay activity in area LIP is linked to the eye movement plan (Snyder et al. 1997; Snyder et al. 1998), the close association between delay activity and responses to auditory stimuli suggests that activity in the auditory stimulus period does contain a substantial intentional component. Therefore, behavioral modulation of responses to auditory stimuli in area LIP may primarily reflect selection of auditory stimuli as targets for eye movements.

Chapter 4 Temporal Features

As shown in previous chapters, the responsiveness of area LIP to auditory stimuli is affected both by auditory-saccade training and by the immediate behavioral context in which auditory stimuli are presented. Compared to visual responses, auditory responses are significantly more dependent on behavioral task, and significantly more predictive of delay or saccade activity. These findings suggest that responses to auditory stimuli in area LIP differ from visual responses in the extent to which they reflect the significance of stimuli as potential saccade targets, rather than the specific sensory parameters of stimuli. In other words, auditory and visual responses in area LIP may represent neural signals at different stages of sensorimotor processing. If so, the two types of responses might be expected to differ in time course. To address this possibility, a novel method for analysis of temporal features in spike train data is proposed. Principled algorithms for smoothing and clustering spike trains are derived from the assumption that spike trains represent randomly scaled inhomogeneous Poisson processes. These algorithms are then used to analyze the temporal features of responses to auditory and visual stimuli in area LIP. The analyses demonstrate that most responses to auditory stimuli in area LIP are gradual in onset, weakly excitatory, and relatively long in duration. In contrast, visual responses can be either excitatory or inhibitory, and the excitatory responses often have a fast transient component. Overall, the results suggest that auditory signals enter area LIP through a much more circuitous route than most visual signals. The findings therefore support the hypothesis that auditory responses in area LIP lie farther along the sensorimotor continuum than visual responses.

4.1 Introduction

Previous chapters of this thesis have demonstrated first, that responses to auditory stimuli emerge in area LIP through auditory-saccade training; and second, that auditory responses after training are stronger in the context of a saccade task than in the context of a fixation task. Compared to visual responses, auditory responses are significantly more dependent on behavioral task, and significantly more predictive of delay or saccade activity. While visual responses in LIP appear to have a substantial sensory component, responses to auditory stimuli seem to depend primarily upon the animal's internal state. These findings suggest that responses to auditory stimuli in area LIP lie farther along the sensorimotor continuum than responses to visual stimuli.

If auditory and visual responses in area LIP do indeed represent neural signals at different stages of sensorimotor processing, the two types of signals might be expected to differ in time course. Auditory responses, for instance, might be more gradual in their development than visual responses. The goal of this chapter, therefore, is to answer the question: Do the temporal features of auditory and visual responses in area LIP differ, and if so, how?

This question cannot be addressed with the analysis methods applied in previous chapters, because those methods were based on mean firing rates across the entire stimulus period. A more detailed analysis of neural activity during the stimulus period is clearly necessary. Thus another issue must be considered at the very outset of this chapter: How should temporal features in spike train data be analyzed?

4.1.1 Subintervals

In principle, the basic analysis strategy used in the previous two chapters — subdivide each trial into multiple intervals, then compare mean firing rates in each interval across stimulus conditions — could be applied on a finer scale within the stimulus period to identify temporal structure in auditory or visual responses. Responses with significantly elevated activity in the early but not the later part of the stimulus

period could be classified as “transient” responses; responses with elevated activity only in the later period could be classified “slow-onset sustained” responses; and so on. However, this approach to analysis of temporal features in spike trains has three serious flaws.

- Unlike the stimulus, delay, and saccade/hold periods defined in the previous chapters, the choice of subintervals in the stimulus period cannot be justified in terms of obvious task-related events. In fact, the definition of relevant subintervals within the stimulus period is completely arbitrary, in the absence of strong prior assumptions regarding the temporal structure expected (for instance, the likely duration of transients).
- Even coarse subdivisions of the stimulus period will create intervals of duration ≈ 100 ms. Estimates of mean firing rate over such small subintervals are extremely noisy. This noise in the firing rate estimates may compound statistical difficulties inherent in repeated application of significance tests across multiple subintervals.
- Finally, and most importantly, the subinterval method does not directly exploit the fact that spike trains are stochastic point processes. Ideally, temporal structure in a spike train should be analyzed through explicit investigation of inhomogeneities in the underlying point process.

4.1.2 Spike density functions

Improving upon the subinterval method, many investigators have analyzed temporal structure in spike train data using methods based upon estimation of spike density functions. All such methods involve conversion of spike trains into smoothed, finely discretized peri-stimulus time histograms. Usually, the first step is to convolve individual spike trains with a fixed or adaptive Gaussian kernel (Richmond et al. 1987; Richmond and Optican 1987; Richmond et al. 1990; McClurkin et al. 1991; Gawne et al. 1991; McClurkin et al. 1991; Eskandar et al. 1992; Tovée et al. 1993; McClurkin

and Optican 1996); alternatively, multiple spike trains collected under the same trial conditions may be averaged, with or without additional smoothing (Middlebrooks et al. 1994; Becker and Krüger 1996; Middlebrooks et al. 1998; Xu et al. 1998). The resulting spike density functions are then discretized into small time bins, with bin-width typically ≈ 5 ms. Binned spike density functions are often represented as points in a T -dimensional vector space, where T is the number of time bins.

Following Richmond and Optican (Richmond and Optican 1987; Optican and Richmond 1987), many investigators have used spike density functions to analyze temporal features and stimulus-related information in spike train data (Richmond and Optican 1990; McClurkin et al. 1991; Gawne et al. 1991; McClurkin et al. 1991; Eskandar et al. 1992; Tovée et al. 1993). The first step in this method is to determine, within the T -dimensional vector representation, the principal components (eigenvectors of the covariance matrix) for all spike density functions from the same neuron. The first few principal components, reconstructed into temporal waveforms, are assumed to capture the dominant temporal features of that neuron's response. Similarities across neurons in the shapes of the first few principal components may then be taken as evidence that certain temporal features are characteristic of neural responses in the brain area under study. Moreover, low-dimensional representations of the spike density functions for a given neuron may be obtained by projecting the T -dimensional representations of the data onto the subspace spanned by the first few principal components. The discriminability of neural responses from different trial conditions may then be quantified through application of information theoretic analysis to these low-dimensional representations, with appropriate corrections for bias due to undersampling (Optican et al. 1991; Tovée et al. 1993; Treves and Panzeri 1995; Golomb et al. 1997). Alternatively, discriminability of trial conditions can be estimated directly from the spike density functions using clustering techniques (Chee-Orts and Optican 1993; McClurkin and Optican 1996), or neural networks may be used either to classify spike density functions by trial condition (Kjaer et al. 1994; Middlebrooks et al. 1994; Heller et al. 1995; Becker and Krüger 1996; Middlebrooks et al. 1998; Xu et al. 1998) or to investigate the nature of the neural code for different

stimulus parameters (Eskandar et al. 1992; McClurkin et al. 1996).

All these spike density function techniques sidestep two pitfalls of the subinterval approach to spike train analysis. First, these methods do not rely upon an arbitrary definition of subintervals (although a width for the time bins and the Gaussian smoothing function must be chosen). Second, these approaches avoid some of the statistical difficulties of the subinterval method, by using principal components decomposition, clustering techniques, or neural networks to reduce the dimensionality of the data and to exclude from analysis those response variations which are most likely to represent noise. However, spike density function techniques do not correct, or even address, the third shortcoming of the subinterval method. Like the subinterval method, these approaches fail to exploit the fact that spike trains are point processes; instead, they employ an ad-hoc smoothing or averaging step to convert spike trains into continuous spike density functions. Moreover, the spike density function approaches introduce some potential drawbacks of their own. By representing spike trains as points in a T -dimensional vector space, these methods make two implicit, and possibly invalid, assumptions: first, that the space of all possible spike trains has a Euclidean geometry; and second, that the sequential nature of time is irrelevant (i.e., that all coordinates in the vector space are equivalent).

4.1.3 Distance metrics

An entirely different approach is exemplified by the recent work of Victor, Purpura, and colleagues (Victor and Purpura 1996; Victor and Purpura 1997; Victor and Purpura 1998; Mechler et al. 1998). Modifying genetic sequence similarity measures for use in spike train analysis, these investigators describe and apply a family of metrics which can be used to calculate distances between spike trains. All the metrics define the distance between two spike trains to be the minimum accumulated cost associated with turning one spike train into the other by repeatedly: (1) adding/deleting a spike, or (2) altering the relative positions of existing spikes. At the optimal setting for the parameter controlling the cost tradeoff between the two steps in the distance

calculation, the mutual information between stimulus and spike trains will be maximized. This optimal cost-parameter setting can then be used to define the scale of stimulus-related temporal structure in the spike trains. Furthermore, once the final pairwise distances have been computed, similarities between spike trains from different trial conditions may be quantified by application of multidimensional scaling and standard clustering techniques.

This metric-space method overcomes many of the problems with the subdivision and spike density function approaches. The method is statistically well-founded, and does not require definition of arbitrary subintervals. Moreover, time between spikes plays a natural role in the distance calculations, and the space of possible spike trains is not assumed to have a Euclidean geometry. But despite these advantages, this method is not appropriate for the purposes of this chapter, because stimulus-locked temporal features of the response are not preserved in the distance metric. Two spike trains at the same distance from a third spike train might have radically different overall response profiles, provided that local temporal features are similar. Thus the distance metric provides no information about the evolution of neural responses over long periods of time relative to stimulus onset. Furthermore, while the metric-space approach does treat spike trains as sets of discrete impulses, the method is not explicit in its representation of the underlying point process.

4.1.4 Hidden Markov models

Recently, several investigators have demonstrated that hidden Markov models (HMMs), originally developed for use in speech recognition tasks, can also be a powerful tool for analyzing spike trains recorded simultaneously from multiple cells (Radons et al. 1994; Abeles et al. 1995; Seidemann et al. 1996; Gat et al. 1997). Each set of simultaneously recorded spike trains is assumed to represent a multivariate Poisson process, in which the underlying rate vector varies over time. The dynamics of the rate variation are modelled as a first-order Markov chain, and standard optimization algorithms are used to estimate the parameters of the hidden Markov model. These

parameters (state transition probabilities, and joint distributions of the spike trains at each hidden state) segment the simultaneously recorded spike trains into a sequence of statistically discriminable hidden states. The hidden states presumably represent distinct modes of brain activity, and have been shown to correlate with behavior (Abeles et al. 1995; Gat et al. 1997).

Both statistically and biologically well-motivated, the HMM approach has none of the problems of the previous models. In contrast to the subinterval method, which imposes a discretization upon the data, the HMM method discovers a natural discretization within the neural activity. Moreover, unlike the spike density function techniques, the HMM method requires no pre-processing of spike trains, and all assumptions regarding the structure of spike trains are explicit in the model. Finally, in contrast to all three previous methods, the HMM approach treats spike trains as stochastic point processes. The HMM method is therefore the most theoretically appealing of the four approaches discussed here. Unfortunately, however, this approach is feasible only when data from a large number of simultaneously recorded cells is available. For analysis of temporal features in spike trains recorded from single units during a limited number of trials, the HMM method is impractical (N. Tishby, personal communication). Therefore, this chapter adopts a novel approach, similar in principle to the HMM method but appropriate for analysis of temporal features in small single-unit data sets.

4.1.5 Randomly scaled inhomogeneous Poisson model

Like the HMM method, the method used in this chapter to analyze temporal features in neural spike train data is derived from a probabilistic model for spike train generation. This generative model is a special case of a doubly stochastic Poisson process — that is, a Poisson process in which the Poisson rate parameter itself is a random variable (Cox 1955; Snyder and Miller 1991). Spike trains collected from the same neuron under identical experimental conditions are assumed to arise from inhomogeneous Poisson processes, whose underlying rates are randomly scaled copies

of a single smooth time-varying function. The choice of this randomly scaled inhomogeneous Poisson model is motivated by recent studies which suggest that slow variations in neural excitability combine with Poisson-like noise to produce the high variability observed in cortical spike trains (Brody 1998; Oram et al. 1998; also see Tomko and Crapper 1974; Rose 1979; Tolhurst et al. 1981; Dean 1981; Arieli et al. 1996). Changes in neural excitability are assumed to result in stochastic scaling of an underlying inhomogeneous Poisson rate profile.

Given this probabilistic generative model, a *temporal feature* can be explicitly defined as the underlying time-varying Poisson rate profile which is scaled from trial to trial to generate a set of spike trains. Furthermore, the model may be used to derive principled algorithms for smoothing and clustering spike trains, which can then be applied to the problem of identifying temporal features common to multiple cells. The probabilistic generative model therefore forms a solid statistical foundation for the analysis of temporal features in spike trains. Like the HMM method, this approach embraces the fact that spike trains are point processes, and provides an explicit probabilistic framework for analysis of neural data. Assumptions about the spike generating process, and additional assumptions regarding parameter values, are clearly laid out in the model, rather than being hidden in the implementation of the algorithm.

Details of the probabilistic generative model, fitting algorithms, clustering techniques, and model selection procedures are provided in the Methods section (and in Sahani 1999). The Results section begins by further motivating the investigation of temporal features in auditory and visual responses, and by testing some of the assumptions of the randomly scaled inhomogeneous Poisson model. Fits of the model to real data are then evaluated, and used to estimate the latencies of auditory and visual responses in area LIP. Finally, clustering algorithms derived from the model are applied to the data, to identify temporal features characteristic of responses to auditory and visual stimuli in area LIP. This analysis shows that the temporal features of auditory and visual responses in LIP do indeed differ. Moreover, the nature of these differences suggests that auditory signals reaching area LIP are much more

highly processed than visual signals. The results also demonstrate that the randomly scaled inhomogeneous Poisson approach is a practical and effective method for analyzing temporal features in spike trains recorded from single neurons. The randomly scaled inhomogeneous Poisson method was developed in collaboration with Maneesh Sahani, and is further discussed in Sahani (1999).

4.2 Methods

4.2.1 Database and data analysis

Results presented in this chapter are based on analysis of 160 unit recordings from LIP neurons in two *Macaca mulatta* monkeys. As explained in Chapter 3, these recordings were collected while the animals performed auditory and visual trials of two different tasks: a memory-saccade task, and a fixation task. See the Methods section of Chapter 3 for a description of the behavioral tasks, and the Methods section of Chapter 2 for details regarding animal care, surgical procedures, experimental setup, recording techniques, and histology.

This chapter focuses on the temporal structure of responses to auditory and visual stimuli, recorded in both the memory-saccade task and the fixation task while the monkeys were fixating straight ahead. The relevant analysis interval in each trial extends for 1500 ms: from 500 ms before stimulus onset, through the 500 ms auditory or visual stimulus presentation period, to 500 ms after stimulus offset. For some analyses of spike train statistics, firing rates or spike counts over two subintervals are considered: the *pre-stimulus* period (the 500 ms interval before stimulus onset), and the *stimulus* period (the 500 ms interval from stimulus onset to stimulus offset).

Auditory stimuli (5–10 kHz, 70 dB SPL noise bursts with 5 ms rise/fall times) or visual stimuli (70 cd/m² red lights covering 0.4°) were presented at one of two possible stimulus locations on each trial: contralateral or ipsilateral to the recording chamber, at positions (−16°, +8°) or (+16°, +8°) relative to the fixation point. Trials of the same task type were run in blocks, but auditory and visual (and contralateral and

ipsilateral) trials were interleaved within each task block. For most of the analyses discussed in this chapter, only responses to stimuli contralateral to the recording chamber are considered.

4.2.2 Probabilistic generative model

As explained in the Introduction section, the algorithms used in this chapter to analyze temporal features in spike train data are derived from a probabilistic generative model for spike train generation. For background information on probabilistic generative models and their use in neural data analysis, see Sahani (1999). The generative model used here is a special case of a doubly stochastic Poisson process, called a randomly scaled inhomogeneous Poisson process (or an inhomogeneous Polya process; see Snyder and Miller 1991). Spike trains collected from the same neuron under identical experimental conditions are each modelled as an inhomogeneous Poisson processes. The underlying rates of these processes are randomly scaled copies of a single smooth time-varying function (the *intensity profile*). The random *scale factor*, by which the intensity profile is multiplied, is drawn independently for each trial from a gamma distribution with unit mean.

Figure 4.1 illustrates the generation of a single spike train according to this model. First, a scale factor is drawn from the cell’s scale factor distribution (Figure 4.1 *A*, gray line); in this example, the chosen scale factor (indicated by the black vertical line) is approximately 1.6. The intensity profile for the cell (Figure 4.1 *B*, gray line) is multiplied by this scale factor to obtain a time-varying Poisson rate parameter for the trial (black line). The spike train (Figure 4.1 *C*, black rasters at top) is a single realization of the inhomogeneous Poisson process characterized by this time-varying rate. Additional spike trains (Figure 4.1 *C*, gray rasters), generated through iteration of this entire procedure, are other instances of the same doubly stochastic Poisson process. Rasters in Figure 4.1 *C* (and in all other raster plots in this chapter) indicate the number of simulated spikes in each 5 ms time bin; the size of each dot is scaled to reflect the number of spikes in the bin.

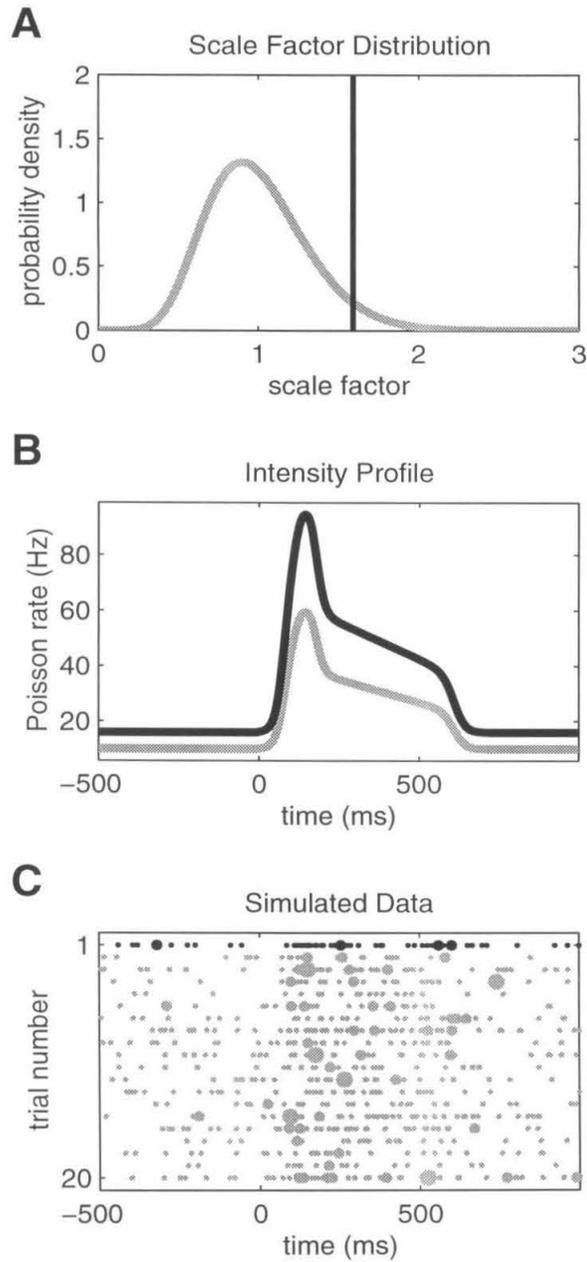


Figure 4.1: The probabilistic generative model which underlies the algorithms for temporal feature analysis. Spike trains are assumed to be realizations of a randomly scaled inhomogeneous Poisson process (see text).

Three considerations underlie the decision to model the scale factor probability density as a gamma distribution with unit mean:

- The multiplicative scale factor is assumed to represent only the overall excitability of a cell on a given trial, not the temporal shape of the cell’s response to a stimulus. Therefore, negative scale factors cannot be allowed, because they would invert the intensity profile. Appropriately, the gamma distribution has no probability density below zero.
- Practical concerns dictate that the scale factor distribution should be as mathematically simple as possible, because equations for the model likelihood and posterior probability must be solved to optimize the parameters of the model. The gamma distribution is a convenient option.
- The scale factor distribution must have unit mean to ensure that the intensity profile represents the mean inhomogeneous Poisson rate across trials. This constraint to unit mean takes a particularly convenient form for the gamma distribution; it reduces the number of parameters in the gamma distribution to one, hereafter termed the *stability parameter*.

As the stability parameter increases, the gamma distribution becomes more tightly distributed around 1. (For reference, the scale factor distribution shown in Figure 4.1 has stability parameter 10; a distribution with stability parameter 100 would have approximately one-third the width.) The stability parameter therefore quantifies the consistency of the neuron’s response from trial to trial, while the intensity profile captures the dominant time-locked temporal feature of the response.

To simplify calculation of the model likelihood and posterior probability, the intensity profile and spike trains are discretized into T bins. A band-limited Gaussian process prior, explained later in this section, is adopted to ensure that the intensity profile will be smooth; as a result, the number of effective parameters is set by the bandwidth of the prior, so the time bins may be made arbitrarily narrow without incurring a complexity penalty. However, computation time does increase with the

number of time bins, so for the analyses presented in this chapter T was chosen to be 300 (i.e., the 1500 ms analysis interval was divided into time bins of width 5 ms each). The model joint probability of observing a single spike train \mathbf{x} and scale factor s , given the intensity profile $\boldsymbol{\lambda}$ and the stability parameter α , is:

$$P(\mathbf{x}, s \mid \boldsymbol{\lambda}, \alpha) = \left(\prod_{t=1}^T \frac{e^{-s\lambda_t} (s\lambda_t)^{x_t}}{x_t!} \right) \left(\frac{\alpha^\alpha}{\Gamma(\alpha)} s^{\alpha-1} e^{-s\alpha} \right) \quad (4.1)$$

The parameters of the model are the discretized intensity profile $\boldsymbol{\lambda}$ and the stability parameter α ; the standard gamma parameters for the scale factor distribution are $(\alpha, 1/\alpha)$. The scale factor s for each spike train is an unobserved (hidden) variable. It is possible to obtain a closed form for the marginal density of the spike trains by integrating over the scale variable. For a set of N spike trains $(\mathbf{x}_1, \dots, \mathbf{x}_N)$ collected from the same cell under identical trial conditions,

$$P(\mathbf{x}_1, \dots, \mathbf{x}_N \mid \boldsymbol{\lambda}, \alpha) = \prod_{n=1}^N \left(\prod_{t=1}^T \frac{\lambda_t^{x_{nt}}}{x_{nt}!} \right) \left(\frac{\Gamma(X_n + \alpha)}{\Gamma(\alpha)} \right) \alpha^\alpha (\Lambda + \alpha)^{-(X_n + \alpha)} \quad (4.2)$$

where the subscript n indicates the spike train, and t indicates the time bin. Λ and X_n are the sums of the elements in the corresponding vectors: $\Lambda = \sum_{t=1}^T \lambda_t$ and $X_n = \sum_{t=1}^T x_{nt}$.

Equation 4.2 is the likelihood function for the parameters $\boldsymbol{\lambda}$ and α , given the observed spike trains $\mathbf{x}_1, \dots, \mathbf{x}_N$. This likelihood must be combined with priors on the intensity profile $\boldsymbol{\lambda}$ and the stability parameter α to obtain the posterior probability function. The prior on α is chosen to be $e^{-1/\alpha}$, which penalizes models with extremely low stability parameters. The prior on $\boldsymbol{\lambda}$ is taken to be a Gaussian process prior with band-limited covariance matrix \mathbf{C} . This band-limited Gaussian process prior ensures that $\boldsymbol{\lambda}$ will tend to be smooth, and is chosen such that the effective temporal resolution of the intensity profile parameter is limited to a scale of about 50 ms (i.e., the cutoff frequency is 10 Hz, which is equivalent to a half-cycle duration of 50 ms). From Bayes' rule, the log posterior is the sum of the log priors, the log likelihood, and a

normalization constant $\log K$ (which absorbs terms independent of $\boldsymbol{\lambda}$ or α):

$$\begin{aligned} \log P(\boldsymbol{\lambda}, \alpha \mid \mathbf{x}_1, \dots, \mathbf{x}_N) &= \log K - \frac{1}{2} \sum_{r=1}^T \sum_{t=1}^T \lambda_r \mathbf{C}_{rt}^{-1} \lambda_t - \frac{1}{\alpha} \\ &+ \sum_{n=1}^N \left\{ \sum_{t=1}^T (x_{nt} \log \lambda_t) - (X_n + \alpha) \log(\Lambda + \alpha) \right. \\ &\quad \left. + \alpha \log \alpha + \log \left(\frac{\Gamma(X_n + \alpha)}{\Gamma(\alpha)} \right) \right\} \end{aligned} \quad (4.3)$$

This log posterior may be maximized by iterative gradient ascent. In practice, it proves useful to express the stability parameter α in the log domain as $\beta = \log \alpha$, and the intensity profile parameter $\boldsymbol{\lambda}$ in the Fourier-transformed log domain as $\boldsymbol{\phi}_\omega = \mathbf{F} \log \lambda_t$. Here \mathbf{F} is a rectangular matrix representation of the discrete Fourier transform operator restricted to the passband of the prior. The logarithmic representation ensures that both λ_t and α remain positive, as is required for a Poisson rate and for the scaling factor; the logarithm also serves to avoid singularities in the gradient at $\lambda_t = 0$ or $\alpha = 0$. Additionally, the Fourier representation for $\boldsymbol{\lambda}$ allows the band-limited prior to be imposed simply by dropping the high-frequency coefficients from the representation. The corresponding model in the time domain would be insurmountably singular. In the Fourier-log domain for $\boldsymbol{\lambda}$ and the log domain for α , the log posterior becomes:

$$\begin{aligned} \log P(\boldsymbol{\phi}, \beta \mid \mathbf{x}_1, \dots, \mathbf{x}_N) &= \log K - \frac{1}{2} e^{\boldsymbol{\phi}^T \mathbf{F}} \mathbf{R} e^{\mathbf{F}^T \boldsymbol{\phi}} - \frac{1}{e^\beta} \\ &+ N \langle \mathbf{x} \rangle^T \mathbf{F}^T \boldsymbol{\phi} - N \langle \langle \mathbf{x} \rangle^T \mathbf{1} + e^\beta \rangle \log(e^{\boldsymbol{\phi}^T \mathbf{F}} \mathbf{1} + e^\beta) \\ &+ N e^\beta \beta + \sum_{n=1}^N \log \left(\frac{\Gamma(X_n + e^\beta)}{\Gamma(e^\beta)} \right) \end{aligned} \quad (4.4)$$

where $\langle \mathbf{x} \rangle$ represents the mean spike train vector over the N different observations, $\mathbf{1}$ is a vector of ones, $\mathbf{R} = \mathbf{F}^T (\mathbf{F} \mathbf{C} \mathbf{F}^T)^{-1} \mathbf{F}$, and \mathbf{T} indicates a transpose. All vectors are column vectors unless transposed, and exponentiation of a vector term is taken to apply element by element.

Gradient ascent is generally most efficient when the initial parameter conditions for the gradient ascent procedure are chosen as follows. The initial intensity profile

λ_0 is set to be flat, with constant magnitude equal to the mean spike count per trial $\langle X \rangle$ divided by the total number of time bins T . Thus λ_0 is the optimal estimate of the *homogeneous* Poisson rate.

$$\lambda_0 = \frac{\langle X \rangle}{T} \mathbf{1} \quad (4.5)$$

The initial stability parameter α_0 is also estimated from the data. Variations in neural excitability across trials will be reflected in the spread of the trial-by-trial spike counts X_n around the mean spike count across trials $\langle X \rangle$. Since the variance of the unit-mean gamma distribution with stability parameter α is equal to $1/\alpha$, the initial stability parameter α_0 may be approximated as the inverse of the variance in the spike count distribution, normalized by the mean spike count across trials. This approximation underestimates the true stability (because it ignores the effects of Poisson variability on the normalized spike count variance), but in practice proves to be a reasonable initial choice.

$$\alpha_0 = \frac{1}{\text{Var}\left(\frac{X_n}{\langle X \rangle}\right)} \quad (4.6)$$

With these initial conditions, the gradient ascent procedure usually converges quickly to a maximum in the log posterior. The parameters λ and α which maximize the log posterior are optimal estimates of the intensity profile and the stability parameter for the set of spike trains $(\mathbf{x}_1, \dots, \mathbf{x}_N)$. By design, these estimates take into account both variability arising from the assumed Poisson nature of the spike trains, and any additional variability arising from changes in neural excitability across trials.

4.2.3 Mixture models

As explained above, Equations 4.1–4.6 may be used to estimate the optimal intensity profile and stability parameter for a set of N spike trains collected from a single cell under identical experimental conditions. The same equations may also be used to estimate the optimal intensity profile and stability parameter for a set of spike trains produced by more than one cell. Figure 4.2 shows a collection of eight different spike train *groups*, each representing a set of spike trains collected from a single simulated

cell. Under the assumptions of the probabilistic generative model, this collection of eight different spike train groups could be modelled as eight randomly scaled inhomogeneous Poisson processes, each characterized by a pair of intensity profile and stability parameters. However, in this particular example, some of the spike train groups have very similar response features; the intensity profiles and stability parameters for those spike train groups are identical. The model which best represents this simulated data set contains only four distinct processes, with the intensity profiles shown above and below the spike train groups (stability parameters are not shown).

Assuming that a data set containing G different spike train groups could be represented by M different intensity profiles and stability parameters (with $M \leq G$), the optimal estimates for the parameters of the M different processes can be obtained by optimizing the log posterior of a *mixture model*. The mixture model consists of M distinct randomly scaled inhomogeneous Poisson processes, along with *mixture probabilities* which capture, for each of the M processes, the probability that the process generated the G groups of spike trains. Thus if a single component of the mixture model (i.e., a single intensity profile and stability parameter) were very likely to have generated half of the G groups of spike trains, then this one component of the model would have mixture probability 0.5. In Figure 4.2, $G = 8$ and $M = 4$; each mixture component represents two out of the eight spike train groups, so the mixture probability for each component is 0.25.

The parameters of an M -component mixture model (M intensity profiles, stability parameters, and mixture probabilities) may be fit to a data set using a variant of the Generalized Expectation-Maximization (GEM) algorithm (Dempster et al. 1977), in which spike trains belonging to the same group (i.e., spike trains collected from the same cell under identical experimental conditions) must all be generated from the same mixture component (see Sahani 1999 for a detailed discussion of this approach). Thus the responsibility of each mixture component for a given spike train is constrained to be the same as its responsibility for all other spike trains from the same group. With this modification to standard GEM in mind, equations for the E-step and the M-step of the algorithm may easily be derived from the log likelihoods

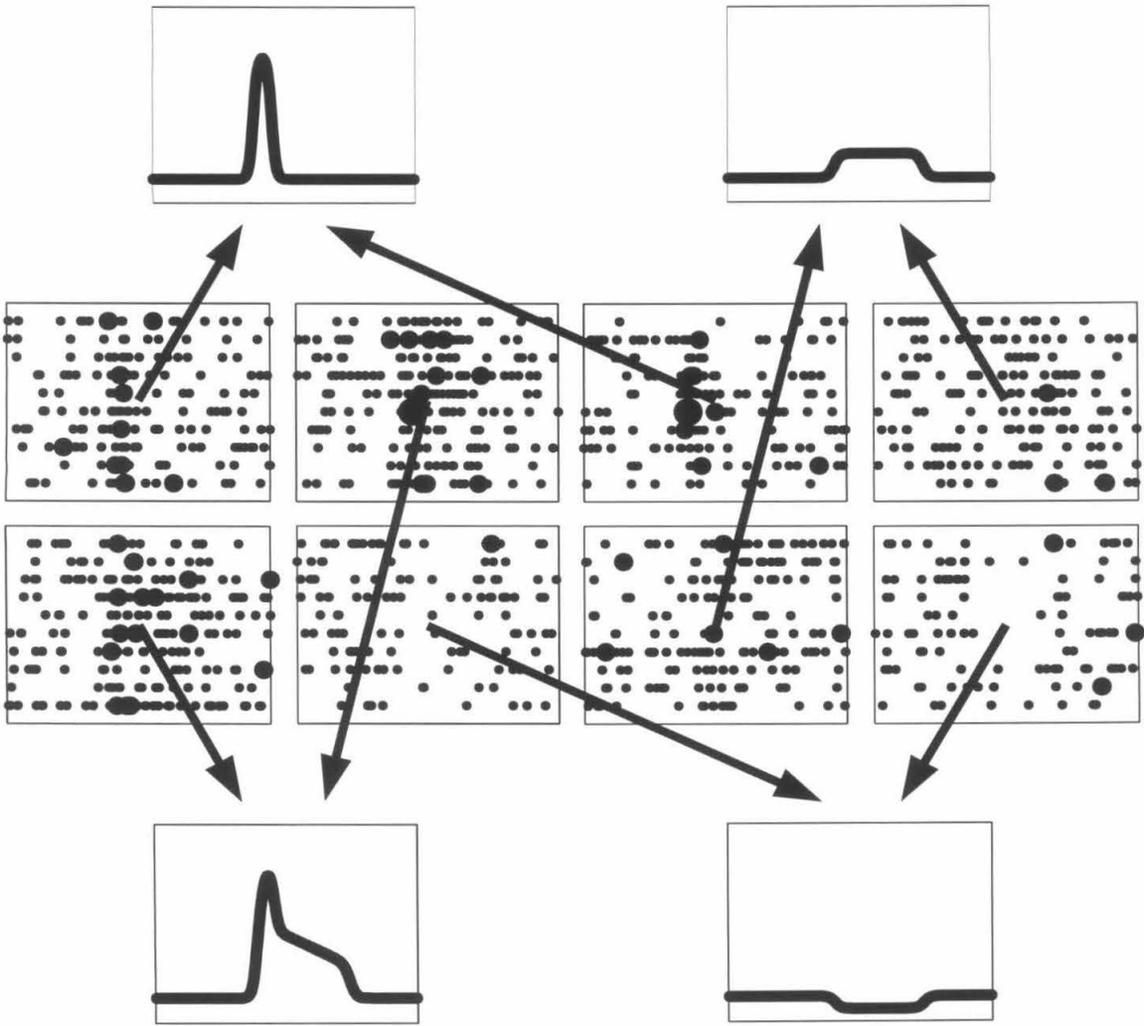


Figure 4.2: A mixture of randomly scaled inhomogeneous Poisson processes. The data set (center) contains spike trains from eight simulated cells, but is best modelled as a mixture of only four randomly scaled inhomogeneous Poisson processes. Intensity profiles for the four processes are shown above and below the spike train groups (stability parameters are not shown). Arrows indicate assignments of spike train groups to processes.

and log posteriors of the model components.

In the E-step, the (unnormalized) responsibility r of model component m for the spike trains in group g is calculated from the probability of group g given model m , weighted by the mixture probability of the model. In the log domain,

$$\begin{aligned} \log r_{m,g} = & \log \pi_m \\ & + \sum_{n=1}^N \left\{ \sum_{t=1}^T (x_{nt} \log \lambda_{m,t}) - (X_n + \alpha_m) \log(\Lambda_m + \alpha_m) \right. \\ & \left. + \alpha_m \log \alpha_m + \log \left(\frac{\Gamma(X_n + \alpha_m)}{\Gamma(\alpha_m)} \right) \right\} \end{aligned} \quad (4.7)$$

where λ_m , α_m , and π_m are the intensity profile, stability parameter, and mixture probability for model m , and the spike trains $\mathbf{x}_1, \dots, \mathbf{x}_N$ all belong to group g . This unnormalized value for $r_{m,g}$ is divided by the total responsibility of all models for group g to obtain the normalized responsibility $\hat{r}_{m,g}$.

$$\hat{r}_{m,g} = \frac{r_{m,g}}{\sum_{m=1}^M r_{m,g}} \quad (4.8)$$

In the M-step, the parameters π_m , λ_m , and α_m are optimized for each model component m . The mixture probability π_m may be calculated directly from the responsibilities:

$$\pi_m = \frac{\sum_{g=1}^G \hat{r}_{m,g}}{\sum_{m=1}^M \sum_{g=1}^G \hat{r}_{m,g}} \quad (4.9)$$

The parameters λ_m and α_m , however, must be optimized by gradient ascent on the weighted posterior probability Q_m :

$$Q_m = \prod_{z=1}^Z (\pi_m P(\lambda_m, \alpha_m \mid \mathbf{x}_1 \dots \mathbf{x}_Z))^{r_{m,z}} \quad (4.10)$$

Here $\mathbf{x}_1, \dots, \mathbf{x}_Z$ represents all the spike trains in the database, and $r_{m,z}$ is the normalized responsibility $\hat{r}_{m,g}$ of model component m for the group g to which spike train \mathbf{x}_z belongs. See Sahani (1999) for further explanation of this step. Logarithmic

transformation and substitution of Equation 4.3 gives:

$$\begin{aligned} \log Q_m &= \log K - \frac{1}{2} \sum_{r=1}^T \sum_{t=1}^T \lambda_{m,r} \mathbf{C}_{rt}^{-1} \lambda_{m,t} - \frac{1}{\alpha_m} \\ &+ \sum_{z=1}^Z r_{m,z} \left\{ \sum_{t=1}^T (x_{zt} \log \lambda_{m,t}) - (X_z + \alpha_m) \log(\Lambda_m + \alpha_m) \right. \\ &\quad \left. + \alpha_m \log \alpha_m + \log \left(\frac{\Gamma(X_z + \alpha_m)}{\Gamma(\alpha_m)} \right) \right\} \end{aligned} \quad (4.11)$$

As before, gradient ascent is easiest if α_m is expressed in the log domain as $\beta_m = \log \alpha_m$, and $\boldsymbol{\lambda}_m$ in the Fourier-transformed log domain as $\boldsymbol{\phi}_{m,\omega} = \mathbf{F} \log \boldsymbol{\lambda}_{m,t}$:

$$\begin{aligned} \log Q_m &= \log K - \frac{1}{2} e^{\boldsymbol{\phi}_m^T \mathbf{F} \mathbf{R}} e^{\mathbf{F}^T \boldsymbol{\phi}_m} - \frac{1}{e^{\beta_m}} \\ &+ \mathbf{r}_m^T \mathbf{X}^T \mathbf{F}^T \boldsymbol{\phi}_m - (\mathbf{r}_m^T \mathbf{X}^T \mathbf{1} + \mathbf{r}_m^T \mathbf{1} e^{\beta_m}) \log(e^{\boldsymbol{\phi}_m^T \mathbf{F} \mathbf{1}} + e^{\beta_m}) \\ &+ \mathbf{r}_m^T \mathbf{1} e^{\beta_m} \beta_m + \sum_{z=1}^Z r_{m,z} \log \left(\frac{\Gamma(X_z + e^{\beta_m})}{\Gamma(e^{\beta_m})} \right) \end{aligned} \quad (4.12)$$

In this equation, \mathbf{X} represents the T -by- Z matrix in which each observed spike train \mathbf{x}_z is a column, and other conventions are the same as in Equations 4.4 and 4.11.

Because the M-step must be performed by gradient ascent, it is computationally much more expensive than the E-step. The Generalized Expectation-Maximization algorithm (in which only a limited number of gradient calculations are performed in each M-step before the responsibilities are updated again with an E-step) is therefore used instead of the standard Expectation-Maximization algorithm (in which the gradient ascent would continue to completion at each step). Thus convergence of the intensity profile and stability parameter for each mixture model component occurs in parallel with convergence of the mixture probabilities.

For the analyses presented in this chapter, convergence under the modified GEM algorithm was judged to be complete either (1) when an iteration of E and M steps produced a change of less than one percent in the combined absolute magnitude of all the intensity profiles and stability parameters in the mixture model, or (2) when the total number of iterations exceeded 100. Initial conditions for the M intensity profile parameters of the mixture model were chosen at random from the optimal

homogeneous Poisson rates for each of the G spike train groups (see Equation 4.5). Initial conditions for the M stability parameters were set to the average of the initial stability estimates for the G groups (see Equation 4.6), and initial mixture probabilities were always $1/M$. Each GEM optimization was restarted 10 times from different initial intensity profiles, and the final mixture model was taken to be the one of the 10 mixture model fits which had the highest log posterior probability. Mixture components with mixture probability less than the floating-point accuracy of the software program used for the fits (MATLAB) were discarded from the final model.

4.2.4 Model selection

If a data set containing G groups of spike trains is known *a priori* to arise from a mixture model containing M different mixture components (with $M \leq G$), then the optimal mixture model is simply the final M -component mixture model obtained through the GEM procedure above. However, for real data sets, the optimal number of mixture components in a data set is itself an unknown variable. Therefore, this number M must also be optimized, using model selection techniques. Mixture models with all possible numbers of components ($1, 2, \dots, G$) must be fit to the data set and then compared.

For all analyses presented in this chapter, model selection was performed using the Bayesian Information Criterion, or BIC (Schwartz 1978). The optimal mixture model was chosen to be the one for which the quantity B_M was maximized:

$$B_M = \left(\sum_{m=1}^M \log P(\mathbf{x}_1, \dots, \mathbf{x}_Z \mid \pi_m, \boldsymbol{\lambda}_m, \alpha_m) \right) - \frac{1}{2}PM \log G \quad (4.13)$$

The first term in this equation is the log likelihood of the entire mixture model under consideration. In the second term, P represents the effective number of parameters per model component, M is the number of mixture components in the model, and G is the total number of spike train groups (the effective number of independent data points). Like other model selection criteria, BIC penalizes the log posterior probability of each mixture model by its complexity; see Sahani (1999) for a review

of model selection techniques.

The effective number of parameters per mixture component is given by:

$$P = \frac{2T}{\tau} + 2 \quad (4.14)$$

where τ is the period (in binwidths) of the cutoff frequency for the band-limited prior on the intensity profile (and as before, T represents the total number of time bins). Thus the number of independent parameters in each intensity profile λ_m is taken to be the length of the intensity profile in half-cycles of the cutoff frequency. Two additional parameters are added for the stability term α_m and the mixture probability π_m . For analyses presented in this chapter, $T = 300$ and $\tau = 20$, so $P = 32$.

Assuming the randomly scaled inhomogeneous Poisson process model is correct, the parameters of the optimal mixture model provide a complete characterization of the data set. The intensity profiles for the M model components capture the dominant temporal features in the neural responses; the stability parameters indicate the consistency of those temporal features across trials and across cells; and the mixture probabilities quantify the relative prominence of the temporal features in the data set. Cells with similar response features may easily be identified as those producing spike trains which are optimally assigned to the same mixture model component, and the relative probabilities of the cell assignments reflect the degree of similarity in the spike trains.

4.3 Results

4.3.1 Population histograms

The initial motivation for this study of temporal features in auditory and visual responses came from examination of population histograms like Figure 4.3. The panels of this figure show, for each of the four different trial conditions, firing rates averaged over all trials recorded from all 160 cells in the database. Mean population activity is plotted as a function of time relative to stimulus onset; the stimulus presentation

period extends for 500 ms. In each plot, divergence between the black line (representing contralateral trials) and the gray line (representing ipsilateral trials) indicates spatial tuning in the population response. Averaged across the population, responses to auditory stimuli seem to emerge gradually in both the memory-saccade task (Figure 4.3 *A*) and the fixation task (Figure 4.3 *C*). In contrast, population responses to visual stimuli are marked by sharp onset transients in both tasks (Figure 4.3 *B* and *D*).

The population averages shown in Figure 4.3 clearly suggest that the responses of LIP neurons to auditory and visual stimuli have different temporal features. However, because data is averaged over thousands of trials to create population histograms, Figure 4.3 provides minimal information about how the temporal features of auditory and visual responses differ at the single-cell level. For example, the fact that the population visual response in Figure 4.3 *B* consists of a large onset transient followed by a smaller sustained component does not necessarily mean that the visual responses of individual cells have a similar profile; this population mean could arise from an average across cells with purely transient responses and cells with purely sustained responses. Likewise, the absence of a transient component in the population auditory response of Figure 4.3 *A* might be deceptive. Cells with transient excitatory auditory responses could conceivably be balanced out, in the population mean, by cells with transient inhibitory responses.

The techniques described in the Methods section were developed in order to make possible quantitative analysis of temporal features at the single-cell level. As previously explained, these analysis methods build on the assumption that spike trains are realizations of doubly stochastic Poisson processes. Spike trains collected from the same neuron under identical experimental conditions are assumed to arise from inhomogeneous Poisson processes, whose underlying rates are randomly scaled copies of a single time-varying rate profile. The random scale factor by which the rate profile is multiplied on each trial is assumed to reflect slow changes in cortical excitability across time.

Many studies have concluded that, to a first approximation, individual cortical

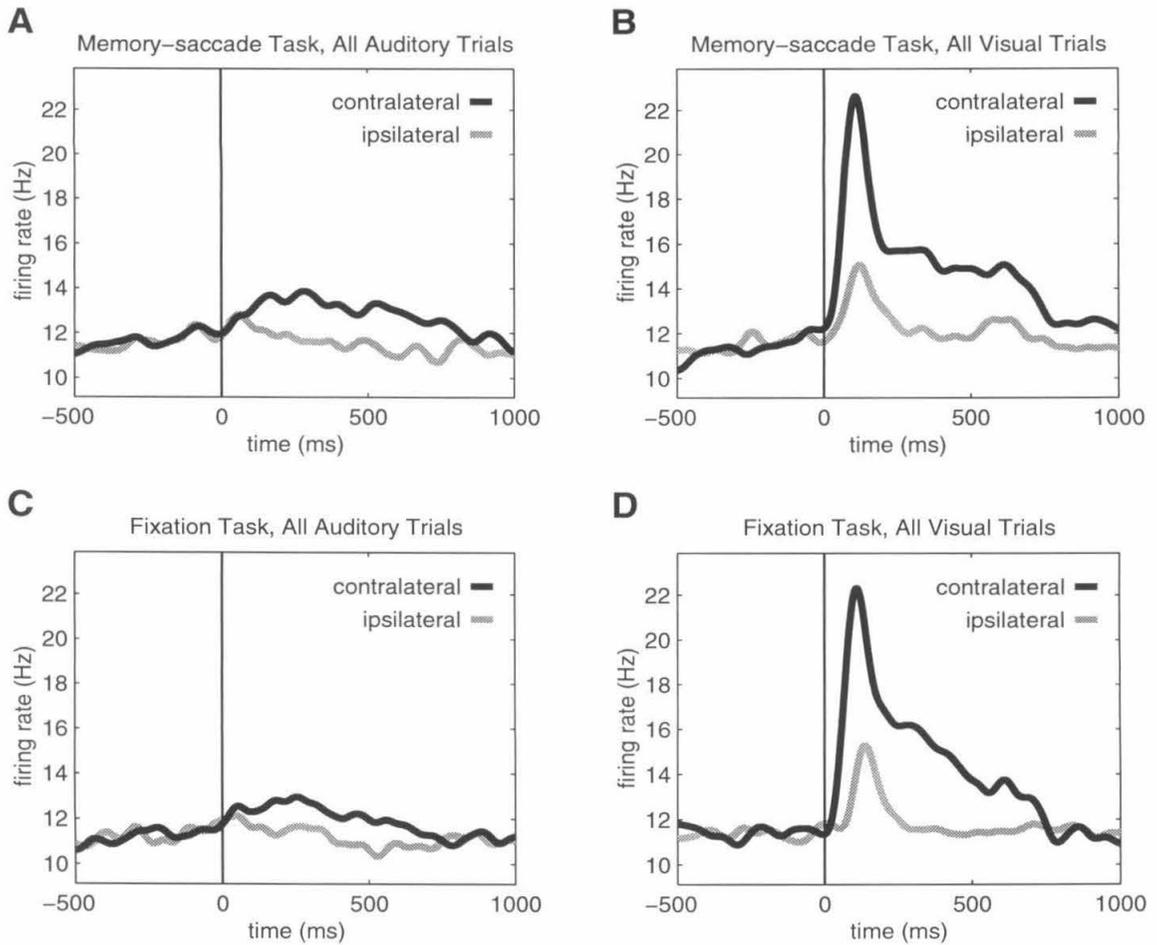


Figure 4.3: Population histograms for each of the four different stimulus/task conditions. Plots show average activity across all trials recorded from all 160 cells in the database, as a function of time relative to stimulus onset (vertical line). Stimulus offset occurs at 500 ms. For each trace, spike trains were discretized into 1 ms bins, averaged across all relevant trials, and then smoothed by convolution with a Gaussian of standard deviation 25 ms.

spike trains are reasonably well modelled as inhomogeneous Poisson processes (Smith and Smith 1965; Moore et al. 1966; Fienberg 1974; Tomko and Crapper 1974; Softky and Koch 1993; Shadlen and Newsome 1996; Bair and Koch 1996). Thus the assumption that spike trains may be approximated as inhomogeneous Poisson processes is not new, and will not be directly analyzed here (although this point is addressed in the Discussion section). The novel aspect of the present model is the assumption that an underlying inhomogeneous Poisson rate profile for each cell and trial condition is randomly scaled across repeated trials. The inclusion of trial-to-trial scaling in the generative model is motivated by recent studies which indicate that variations in cortical excitability across trials can account for many statistical features of cortical spike trains (Brody 1998; Oram et al. 1998). The random scaling component of the generative model may also be justified in part by analysis of spike train statistics in the current database.

4.3.2 Variability across trials

Variability in cortical spike trains is typically characterized in terms of either interspike interval statistics or spike count statistics. In practice, analysis of interspike intervals is problematic, since inhomogeneities in the firing rate over time within each trial (e.g., stimulus-evoked changes in firing) will skew the interval distributions. Previous investigators have attempted to deal with this problem either by analyzing data only over a period in which the mean firing rate of the neuron is relatively constant (reviewed in Moore et al. 1966; also see Shadlen and Newsome 1996), or by using statistical methods to compensate for inhomogeneities in the firing rate (Perkel et al. 1967; Softky and Koch 1993; Holt et al. 1996). When the topic of interest is specifically variability in firing rate across trials (rather than variability in spike timing within trials), these problems may be completely avoided by considering spike count statistics instead. Any fixed interval of an inhomogeneous Poisson process will have the same counting statistics as a homogeneous Poisson process with rate equal to the mean inhomogeneous rate over the interval. Variations in firing rate within each

trial will not complicate analysis of spike counts, provided that the inhomogeneities are time-locked with respect to the analysis interval. Therefore, in this part of the Results section, trial-to-trial variability is analyzed in terms of spike counts and firing rates alone, over intervals fixed with respect to stimulus onset.

Many investigators have noted an apparent power-law relationship between spike count variance and spike count mean, when responses from individual cortical cells are compared across different trial conditions (Tolhurst et al. 1981; Dean 1981; Tolhurst et al. 1983; Gershon et al. 1998). A very similar power-law relationship has been shown to emerge when spike count statistics are compared across different cells for the same trial condition (Vogels et al. 1989). In the present study, the number of different trial conditions was rather small; therefore, the relationship between spike count variance and spike count mean is shown in Figure 4.4 across cells for each stimulus/task condition, rather than across stimulus/task conditions for each cell. Spike count statistics are calculated over the 500 ms stimulus presentation period, for all contralateral trials recorded from each of the 160 neurons in the database. The relationship between spike count variance and spike count mean is nearly linear on a log-log plot, indicating that the well-known power-law relationship holds for this data set. The parameters of the power-law relationship, determined by linear regression in the log domain, are shown for each stimulus/task condition (along with the r -statistic for the regression).

If each cortical cell represented a single inhomogeneous Poisson process, then spike count variance across repeated trials would equal spike count mean for each cell; linear regression in the log domain would produce unity values for both the proportionality coefficient (intercept of the regression line at mean spike count = 1) and the power coefficient (slope of the regression line). In fact, for the four stimulus/task conditions shown in Figure 4.4 *A–D*, the proportionality coefficients are all significantly greater than one ($p < 0.05$). Thus spike count variance significantly exceeds spike count mean across repeated trials. The power coefficients are also greater than one in all four plots, but these values are not significantly different from the expected unity value for Poisson processes (except in Figure 4.4 *C*). Overall, these results are in

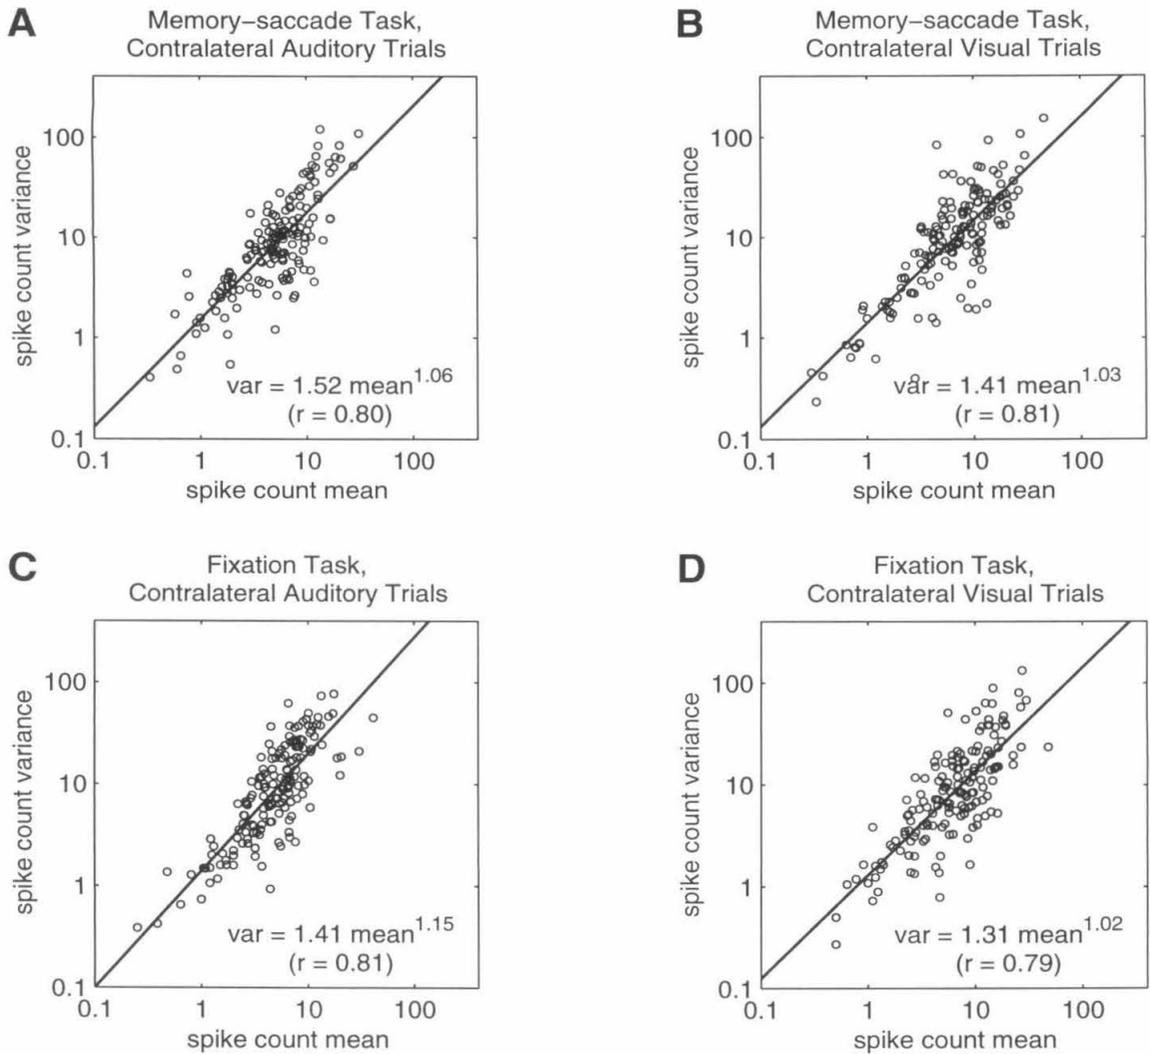


Figure 4.4: The relationship between spike count variance and spike count mean, for each of the four stimulus/task conditions and for all 160 cells in the database. Spike count statistics are calculated for the 500 ms stimulus presentation interval across repeated contralateral trials. The solid line in each plot is the regression line; regression parameters are shown in the power-law relationship. The r -statistic for the log-domain linear fit is indicated in parentheses. The 95% confidence intervals for the proportionality and power coefficients are: *A* [1.22 1.89], [0.94 1.19]; *B* [1.12 1.78], [0.91 1.14]; *C* [1.12 1.77], [1.02 1.28]; *D* [1.02 1.67], [0.89 1.14].

good agreement with those of many previous studies (Tolhurst et al. 1981; Dean 1981; Tolhurst et al. 1983; Vogels et al. 1989; Snowden et al. 1992; Softky and Koch 1993; Gershon et al. 1998; Shadlen and Newsome 1998; but see Gur et al. 1997). Under a variety of different experimental conditions, variance-mean relationships for cortical spike counts are reported to have power coefficients near one, and proportionality coefficients greater than or equal to one. Thus cortical spike trains, including those under consideration in the present study, are often more variable than would be expected from simple Poisson processes.

As many investigators have pointed out, this high variability may arise from non-stationarities in firing rate across repeated trials (Henry et al. 1973; Tomko and Crapper 1974; Rose 1979; Tolhurst et al. 1981; Dean 1981; Tolhurst et al. 1983). More specifically, cortical spike trains might be Poisson processes for which the overall rate fluctuates from trial to trial (i.e., doubly stochastic Poisson processes). One possible form which this hypothesized trial-to-trial fluctuation might take is a *multiplicative* scaling of the overall inhomogeneous Poisson rate profile. The generative model used in this chapter assumes this type of scaling (see Methods). However, an equally simple (and, *a priori*, equally possible) form of trial-to-trial fluctuation would be an *additive* scaling of the overall response. These two hypotheses regarding variation in response from trial to trial may be compared experimentally, because they make different predictions about the relationship between background and stimulus-period firing rates across trials.

Suppose R_n^{back} and R_n^{stim} are firing rates recorded during the pre-stimulus and stimulus periods of trial n , and assume that these rates are scaled by a factor c_n . The multiplicative scaling hypothesis predicts:

$$R_n^{back} = c_n \langle R^{back} \rangle \quad (4.15)$$

$$R_n^{stim} = c_n \langle R^{stim} \rangle \quad (4.16)$$

$$R_n^{stim} - R_n^{back} = c_n (\langle R^{back} \rangle - \langle R^{stim} \rangle) \quad (4.17)$$

$$R_n^{both} = \frac{c_n}{2} (\langle R^{back} \rangle + \langle R^{stim} \rangle) \quad (4.18)$$

where R_n^{both} is the mean rate over both intervals, and $\langle \cdot \rangle$ represents a mean over all trials $1 \dots N$. From these equations, it may be seen that:

$$R_n^{stim} - R_n^{back} = \left(2 - \frac{4 \langle R^{back} \rangle}{\langle R^{stim} \rangle + \langle R^{back} \rangle} \right) R_n^{both} \quad (4.19)$$

Thus the difference in firing rates between the stimulus and pre-stimulus periods depends linearly on the average firing rate over both periods. Linear regression of $R_n^{stim} - R_n^{back}$ against R_n^{both} should produce a line of slope $\left(2 - \frac{4 \langle R^{back} \rangle}{\langle R^{stim} \rangle + \langle R^{back} \rangle} \right)$.

In contrast, the additive scaling hypothesis predicts:

$$R_n^{back} = c_n + \langle R^{back} \rangle \quad (4.20)$$

$$R_n^{stim} = c_n + \langle R^{stim} \rangle \quad (4.21)$$

$$R_n^{stim} - R_n^{back} = \langle R^{back} \rangle - \langle R^{stim} \rangle \quad (4.22)$$

$$R_n^{both} = \frac{1}{2} (2c_n + \langle R^{back} \rangle + \langle R^{stim} \rangle) \quad (4.23)$$

From Equation 4.22, it may be seen that the difference in firing rates between the stimulus and pre-stimulus periods does not depend on the average firing rate over both periods. Therefore, according to this hypothesis, linear regression of $R_n^{stim} - R_n^{back}$ against R_n^{both} should produce a line of slope zero.

Obviously, the two forms of scaling cannot be distinguished when there is no difference in mean firing rate between the stimulus and pre-stimulus periods (that is, when $\langle R^{stim} \rangle - \langle R^{back} \rangle = 0$), because in that case the predicted slope of the regression line would be zero under both hypotheses. Likewise, if the overall firing rate (R_n^{both}) does not vary significantly across trials, then any linear dependence of $R_n^{stim} - R_n^{back}$ on R_n^{both} will be difficult to detect. Figure 4.5 displays data from four cells and trial conditions for which these potential problems are minimized. These four cells and contralateral trial conditions were selected for inclusion in this plot based only on the following two criteria: (1) the difference in firing rate between the stimulus and pre-stimulus periods is significant at the 0.001 level, and (2) the variance in overall firing rate across trials is extremely high (higher than in all other data sets meeting

the first criterion).

In each panel of Figure 4.5, the difference in firing rate between the stimulus and pre-stimulus periods ($R_n^{stim} - R_n^{back}$) is plotted against the overall firing rate for both periods (R_n^{both}) for each trial. The solid line is the linear regression to the data points; the dashed line is the predicted slope under the multiplicative scaling hypothesis; and the dotted line is the predicted slope under the additive scaling hypothesis. Shaded areas indicate the 95% confidence intervals on the slope of the true regression line. For the cells and trial conditions shown in Figure 4.5 *A* and *B*, the scaling is clearly multiplicative; the 95% confidence intervals on the slope of the regression line include the slope predicted under the multiplicative scaling hypothesis but not the zero slope predicted under the additive scaling hypothesis. In contrast, the data shown in Figure 4.5 *C* better support the additive scaling hypothesis; the confidence intervals include the slope of zero (additive scaling), but not the slope which would be expected if scaling were multiplicative. For the cell displayed in Figure 4.5 *D*, the two hypotheses cannot be distinguished, because the confidence intervals include both predicted slopes.

Across the entire database of contralateral trials (160 cells, 4 different stimulus/task conditions per cell), results of the test illustrated in Figure 4.5 are inconclusive for 77% of the data sets (confidence intervals on the regression line include slopes predicted under both hypotheses). As previously explained, there are many situations in which it would not be possible to distinguish between the two hypotheses; therefore, this result is not surprising. More interesting are the remaining data sets: 14% show evidence for multiplicative scaling (confidence intervals include only the predicted slope under the multiplicative scaling hypothesis), while only 5% favor additive scaling (confidence intervals include only the predicted slope under the additive scaling hypothesis). Another 3% provide no evidence for either hypothesis (confidence intervals include neither predicted slope). As a final test of the two scaling hypotheses, the sums of squared deviations of the data from the lines predicted under each of the two hypotheses were computed, data set by data set. These sums of squared deviations were then compared to determine which of the two scaling hypotheses better

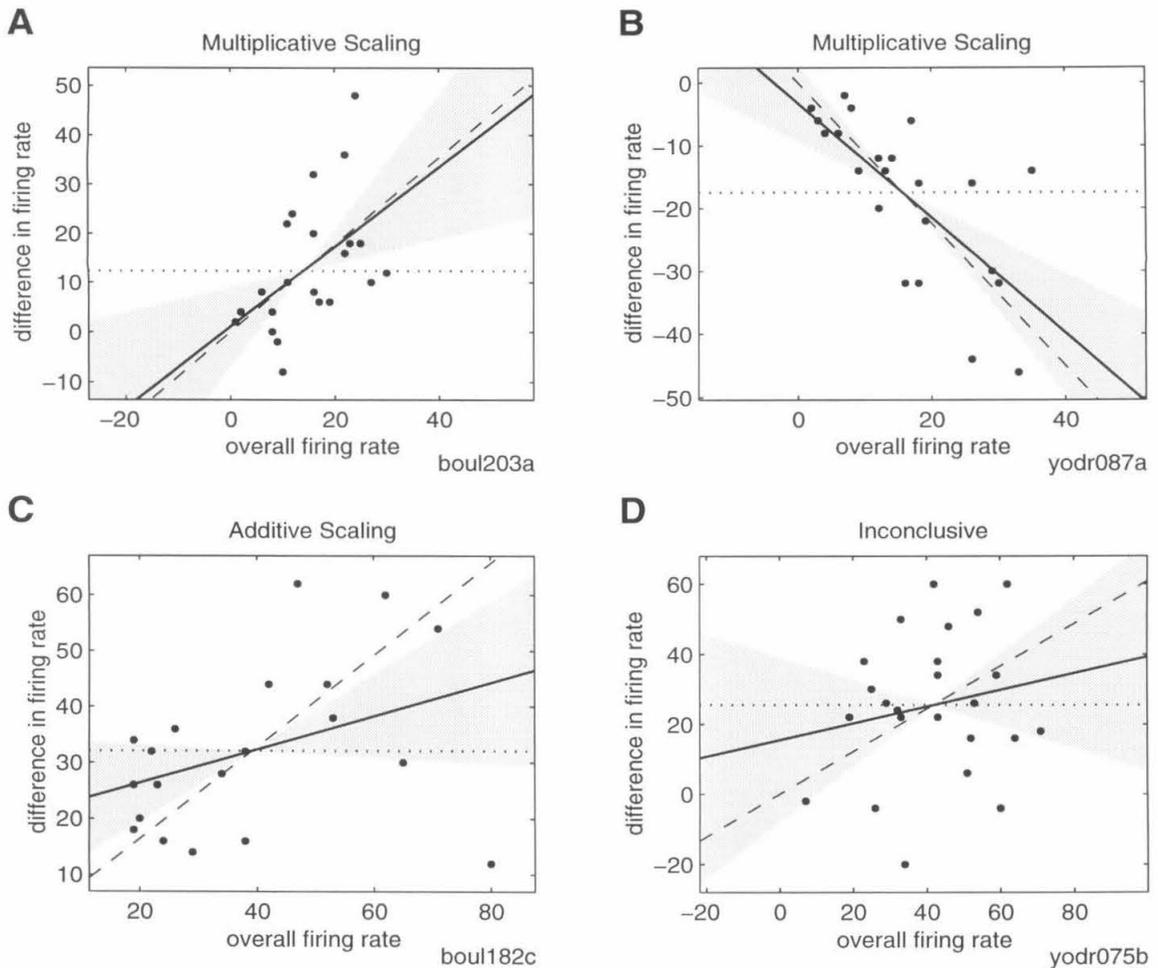


Figure 4.5: A comparison of the multiplicative and additive scaling hypotheses. Data sets are taken from different cells and stimulus/task conditions, and were chosen for inclusion in this figure based on criteria which are unbiased with respect to the two hypotheses under consideration (see text). Each panel shows, for each contralateral trial, the difference in firing rate between the stimulus and pre-stimulus periods plotted against the overall firing rate for both periods. The solid line represents the actual linear regression to the data points; the dashed line is the predicted slope under the multiplicative scaling hypothesis; and the dotted line is the predicted slope under the additive scaling hypothesis. Shaded areas indicate 95% confidence intervals on the slope of the regression line. Type and number of trials in each of the four different data sets: *A* auditory memory-saccade task, 24 trials; *B* visual memory-saccade task, 25 trials; *C* visual fixation task, 20 trials; *D* visual memory-saccade task, 24 trials.

explained each data set. According to this final test, the multiplicative scaling hypothesis provides a significantly better overall fit to the data than the additive scaling hypothesis (binomial test, $p < 0.001$). Thus the weight of the experimental evidence favors the multiplicative scaling hypothesis. The use of multiplicative rather than additive scaling in the probabilistic model for spike train generation (see Methods) is therefore justified.

4.3.3 Goodness-of-fit

The validity of the probabilistic generative model may also be assessed directly, by comparing observed (real) spike trains to simulated spike trains generated from the model. Given a set of observed spike trains $\mathbf{x}_1^{obs} \dots \mathbf{x}_N^{obs}$, a goodness-of-fit test can be conducted as follows.

1. Find the intensity profile $\boldsymbol{\lambda}^{obs}$ and stability parameter α^{obs} which maximize Equation 4.3. This intensity profile and stability parameter define the optimal generative model for the spike trains $\mathbf{x}_1^{obs} \dots \mathbf{x}_N^{obs}$.
2. Calculate $P(\mathbf{x}_1^{obs}, \dots, \mathbf{x}_N^{obs} \mid \boldsymbol{\lambda}^{obs}, \alpha^{obs})$, the likelihood of obtaining the observed spike trains $\mathbf{x}_1^{obs} \dots \mathbf{x}_N^{obs}$ given the optimal generative model described by $\boldsymbol{\lambda}^{obs}$ and α^{obs} (Equation 4.2).
3. Generate a set of simulated spike trains $\mathbf{x}_1^{sim} \dots \mathbf{x}_N^{sim}$ from this optimal generative model, through the procedure illustrated in Figure 4.1.
4. Re-fit the generative model to the simulated data, to find the intensity profile $\boldsymbol{\lambda}^{sim}$ and stability parameter α^{sim} which maximize Equation 4.3 for the simulated spike trains.
5. Calculate $P(\mathbf{x}_1^{sim}, \dots, \mathbf{x}_N^{sim} \mid \boldsymbol{\lambda}^{sim}, \alpha^{sim})$, the likelihood of obtaining the simulated spike trains $\mathbf{x}_1^{sim} \dots \mathbf{x}_N^{sim}$ given the optimal generative model described by $\boldsymbol{\lambda}^{sim}$ and α^{sim} (Equation 4.2).

6. Compare $P(\mathbf{x}_1^{obs}, \dots, \mathbf{x}_N^{obs} \mid \boldsymbol{\lambda}^{obs}, \alpha^{obs})$, the likelihood of the observed spike trains under their optimal generative model, to $P(\mathbf{x}_1^{sim}, \dots, \mathbf{x}_N^{sim} \mid \boldsymbol{\lambda}^{sim}, \alpha^{sim})$, the likelihood of the simulated spike trains under their own optimal generative model. If the probabilistic generative model provides good fit to real neural data, then the likelihoods of observed and simulated spike trains should be similar on average (over many simulated data sets).

It is necessary to re-fit the generative model to the simulated spike trains because the optimal parameters for the generative model fit to the observed spike trains cannot be used to evaluate the likelihood of the simulated spike trains without introducing a bias into the goodness-of-fit estimate. This bias emerges because the generative model fit to a finite set of spike trains will inevitably over-emphasize the random particularities of the spike trains in that data set. Optimizing the model parameters for the simulated spike trains minimizes any bias due to over-fitting. See Sahani (1999) for a more detailed justification of this approach.

The procedure outlined above was repeated for each cell using 100 simulated data sets, to obtain an expected distribution of likelihoods for the observed spike trains. The fit of the generative model to the observed data for each cell was then quantified as the percentile of the simulated likelihood distribution within which the likelihood of the observed data fell. Percentiles below 50 indicate that the likelihood of the observed spike trains was less than the likelihoods of most of the simulated spike trains.

The results of this goodness-of-fit test are illustrated in Figure 4.6, for four cells with very different response profiles. Each panel displays a set of spike trains ($\mathbf{x}_1^{obs}, \dots, \mathbf{x}_N^{obs}$), recorded from a single cell during repeated contralateral trials of the same stimulus/task condition. Because the spike trains are discretized into 5 ms bins, each raster dot can represent multiple spikes; the size of each dot is scaled to reflect the number of spikes falling within the corresponding 5 ms bin. Also shown on each plot is the intensity profile ($\boldsymbol{\lambda}^{obs}$) for the optimal generative model fit to the observed spike trains. The value of the stability parameter (α^{obs}) for each of the four models is noted in the figure legend. As indicated in the panel titles, the fit of the generative

model to all four data sets is acceptable; the likelihood of the observed spike trains under the optimal generative model for each cell falls well within the 5–95 percentile range of simulated data likelihoods.

Figure 4.7 displays, for each of the four cells in Figure 4.6, a single one of the 100 simulations used to estimate the expected likelihood distribution. Rasters in each plot represent simulated spike trains $(\mathbf{x}_1^{sim}, \dots, \mathbf{x}_N^{sim})$, generated from the optimal generative model fit to the real neural data shown in the corresponding panel of Figure 4.6. Superimposed on the simulated data is the intensity profile (λ^{sim}) for the generative model obtained by re-fitting the model parameters to the simulated data. The intensity profiles fit to the real spike trains in Figure 4.6 are roughly similar to the intensity profiles fit to the simulated spike trains in Figure 4.7. However, there are clearly differences in the fine temporal structure of the original and re-fit intensity profiles; these differences arise primarily through over-fitting.

Within each stimulus/task condition, the goodness-of-fit test illustrated in Figures 4.6 and 4.7 was applied to all contralateral trials recorded from cells judged to have a spatially tuned response to the stimulus. In other words, the response criterion outlined in Chapter 2 and Chapter 3 (significant difference between contralateral and ipsilateral firing rates over the stimulus period) was used to select cells and trial conditions for the goodness-of-fit analysis. This restriction to spatially tuned responses was imposed in order to focus the analysis on those cells and stimulus/task conditions with clear stimulus-evoked activity, so that the overall results of the goodness-of-fit tests would not be inflated by good fits to data from unresponsive cells. Moreover, as explained in the previous chapters, only cells with spatially tuned activity are considered because spatially untuned responses cannot be distinguished from general arousal effects. For the vast majority of cells with spatially tuned responses, stimulus-related temporal structure was evident only in spike trains recorded during contralateral trials; therefore, only contralateral trials are analyzed.

Figure 4.8 summarizes the results of the goodness-of-fit tests for all cells with spatially tuned stimulus-period activity. Each panel shows a histogram of model fit percentiles for one stimulus/task condition; N is the total number of cells with

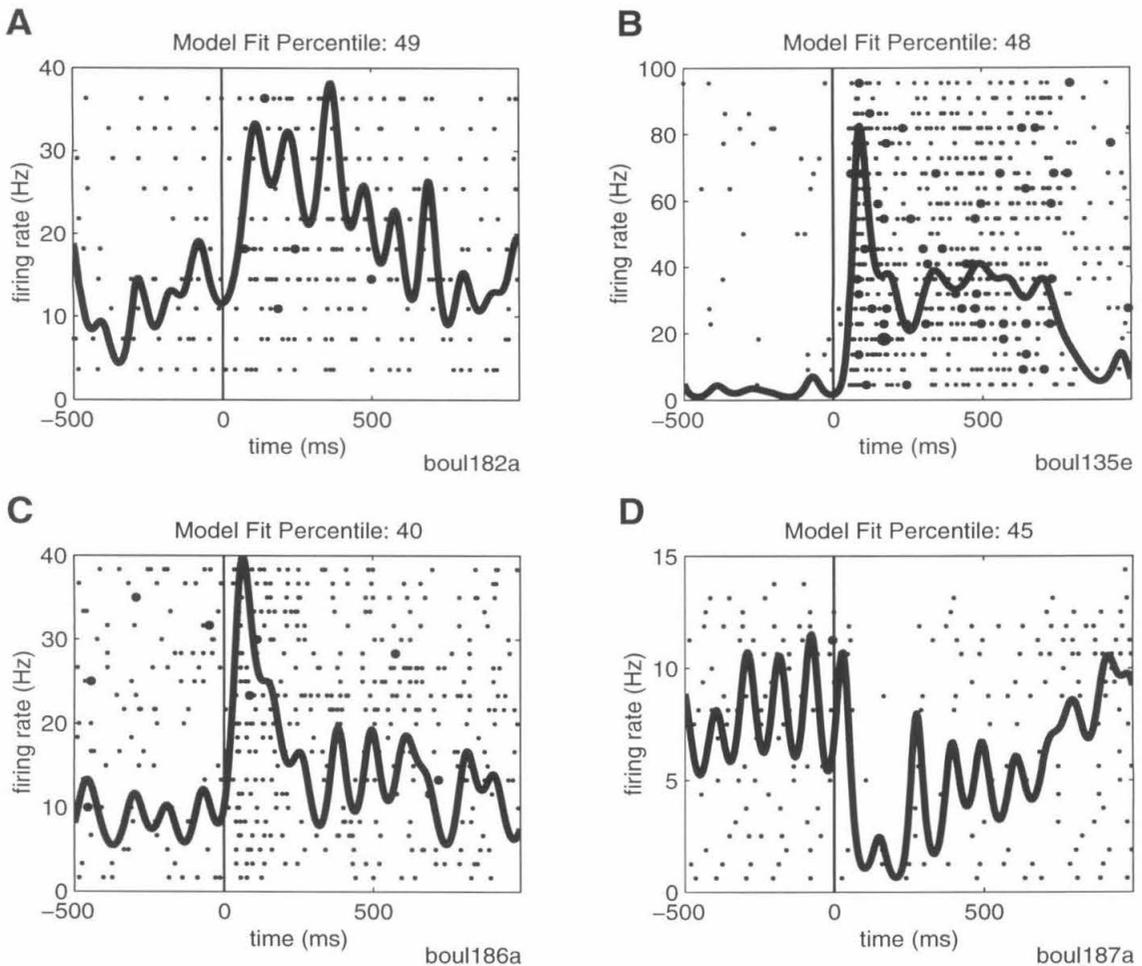


Figure 4.6: Fits of the generative model to real spike train data, for four very different types of stimulus-evoked activity. Panels show spike trains recorded from four different cells, during contralateral trials of the auditory memory-saccade task (*A*), visual memory-saccade task (*B*), auditory fixation task (*C*), or visual fixation task (*D*). Responses shown are not intended to be representative of particular trial conditions; cells were chosen for this plot only in order to illustrate the range of observed stimulus-evoked activity. Spike trains in each panel are discretized into 5 ms bins, and the size of each raster dot is scaled to reflect the number of spikes in the corresponding bin. The thin vertical line in each panel marks the time of stimulus onset; the stimulus presentation interval extends for 500 ms. The thick wavy line in each plot is the intensity profile for the optimal generative model fit to the data set. Note that the vertical scale differs from plot to plot. Stability parameters: *A* 139.4; *B* 49.3; *C* 55.3; *D* 6.2.

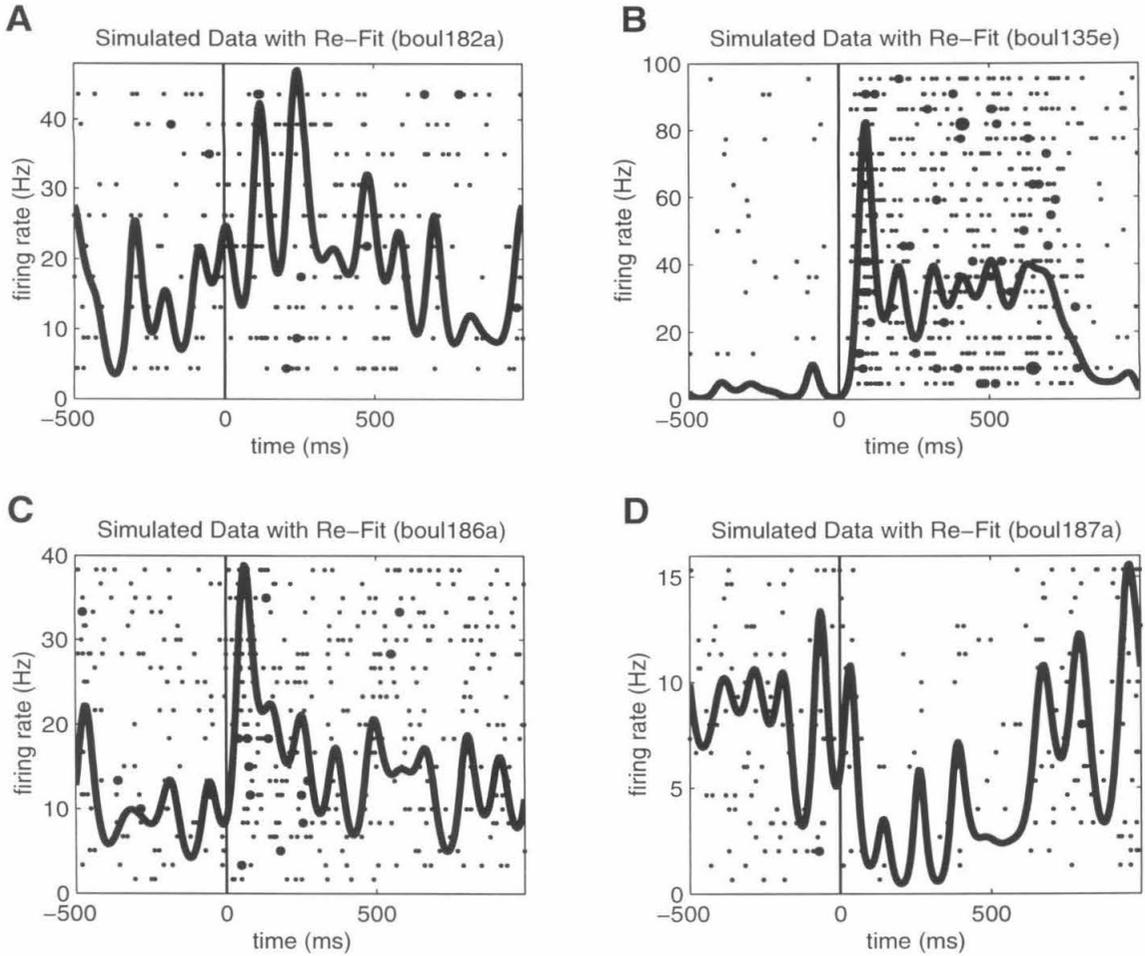


Figure 4.7: Simulated spike trains with re-optimized intensity profiles. Simulated spike trains were generated from models fit to real spike trains in the corresponding panels of Figure 4.6. Conventions are the same as in that figure, but here all data is simulated and intensity profiles represent re-fits to simulated data. Stability parameters for fits to simulated data: *A* 3.2×10^7 ; *B* 53.1; *C* 74.2; *D* 24.5.

spatially tuned stimulus-period activity in that stimulus/task condition. If the assumptions underlying the generative model were correct, then the likelihoods of real data and simulated data would be identically distributed for each cell, and the histograms in each summary plot would be uniform across the 1–100 percentile range. Obviously, the actual histograms are not uniform; for almost every cell analyzed, the model fit percentile is within the 20–60 percentile range. This clumping of the model fit percentiles suggests the likelihoods of simulated data are more widely distributed, on average, than the likelihoods of real data. Therefore, there are clearly some inaccuracies in the assumptions underlying the generative model; spike trains cannot be perfectly modelled as randomly scaled inhomogeneous Poisson processes. This point will be further addressed in the Discussion section. The generative model is, however, good enough that real spike trains cannot be distinguished from simulated spike trains using this goodness-of-fit test. The likelihood of the observed spike trains fell within the expected range of likelihoods for simulated spike trains for every cell except one (Figure 4.8 *D*, one cell at percentile 0), and even for this one cell, the observed likelihood is just at the border of the expected likelihood range (not shown). Overall, then, the fit of the model is acceptable for all the cells tested. Therefore, it seems reasonable to include all cells with spatially tuned stimulus-period activity in further analyses.

4.3.4 Response latencies

Once the generative model has been fit to a set of spike trains from one cell and trial condition, the latency of the neural response may be estimated from the optimal intensity profile. Suppose the mean value of the intensity profile over the pre-stimulus (background) period is μ^{back} , and the standard deviation of the intensity profile over that period is σ^{back} . The response latency may then be defined as the first time after stimulus onset when the absolute value of the intensity profile exceeds $\mu^{back} + 3\sigma^{back}$. This calculation gives a conservative estimate of the response latency. If the absolute value of the intensity profile fails to exceed $\mu^{back} + 3\sigma^{back}$ during the stimulus presen-

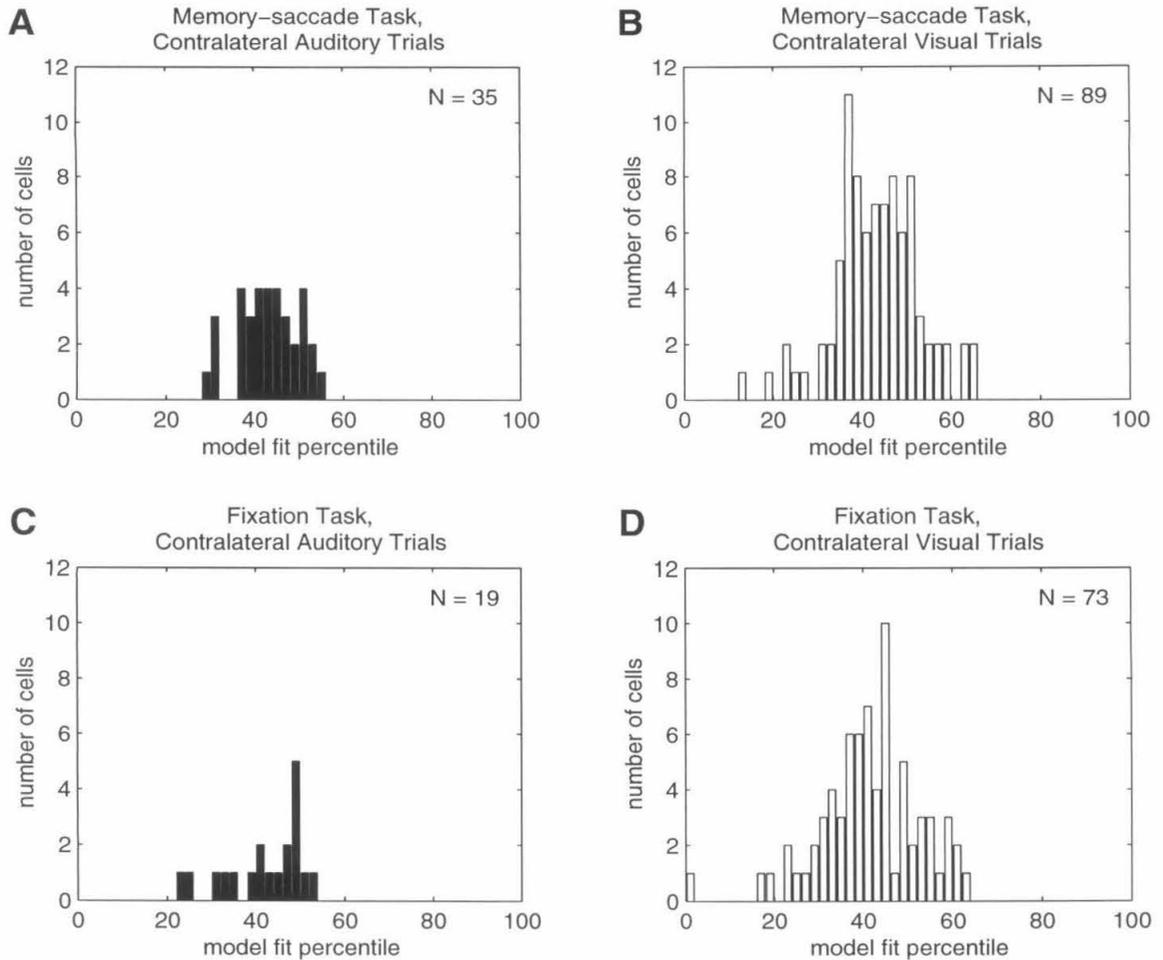


Figure 4.8: Model fit percentile histograms, for each stimulus/task condition. Goodness-of-fit test results are summarized in each panel for contralateral trials from all cells with spatially tuned stimulus-period activity in the indicated stimulus/task condition. N is the total number of cells included in each histogram, and all histogram bins are two percentiles wide.

tation interval, then the cell is judged to have an indeterminate response latency, and is excluded from further analysis.

Figure 4.9 displays spike trains and optimal intensity profiles for contralateral trials from four different cells and stimulus/task conditions. The first thin vertical line in each plot marks the time of stimulus onset; the second thin vertical line indicates the estimated latency of the neural response. Data in Figure 4.9 *A* and *C* are taken from auditory memory-saccade trials and auditory fixation trials, respectively; data in Figure 4.9 *B* and *D* are taken from visual memory-saccade trials and visual fixation trials. These examples demonstrate that the latencies estimated by the procedure described above are reasonably close to those which would be chosen by a human observer. Moreover, although it should be noted that the spike trains shown are from four different cells, the plots illustrate a general observation: auditory responses tend to have longer latencies than visual responses.

This observation is quantified in Figure 4.10. Each histogram shows the distribution of response latencies for a different stimulus/task condition. Response latencies included in each plot were estimated for contralateral trials recorded from cells with spatially tuned stimulus-period activity in the appropriate stimulus/task condition. (The total number of cells in each histogram is smaller than the total number of cells in the corresponding histogram of Figure 4.8, because some cells did not have a clearly identifiable response latency and were therefore excluded from consideration.) Overall, auditory response latencies are significantly longer than visual response latencies in both the memory-saccade task (Figure 4.10 *A* vs. *B*; Mann-Whitney test, $p < 0.01$) and the fixation task (Figure 4.10 *C* vs. *D*; $p < 0.05$). However, auditory response latencies do not differ significantly between the two tasks (Figure 4.10 *A* vs. *C*; Mann-Whitney test n.s.). Visual response latencies, while slightly longer in the fixation task than in the memory-saccade task, also do not differ significantly between the two tasks (Figure 4.10 *B* vs. *D*; n.s.). Therefore, across the entire database, response latencies are dependent on stimulus modality but not on behavioral task. This conclusion still holds when only latencies less than 400 ms are considered (so that the one apparent outlier at latency 425 ms in Figure 4.10 *A* is excluded).

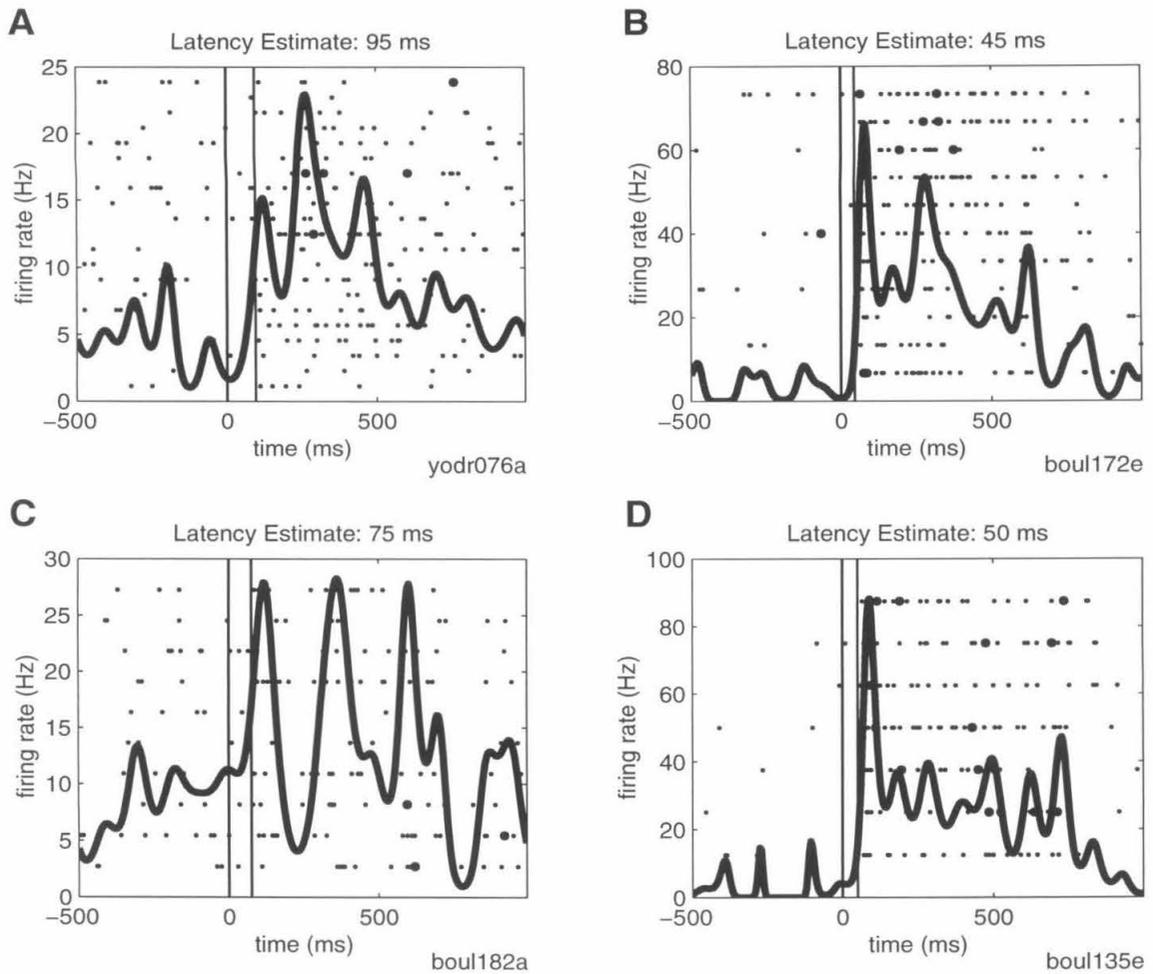


Figure 4.9: Estimated response latencies for four different cells. Panels show examples of spike trains recorded during contralateral trials of the auditory memory-saccade task (*A*), visual memory-saccade task (*B*), auditory fixation task (*C*), or visual fixation task (*D*). The first thin vertical line in each plot marks the time of stimulus onset; the second thin vertical line indicates the estimated response latency. Other conventions are as in Figure 4.6.

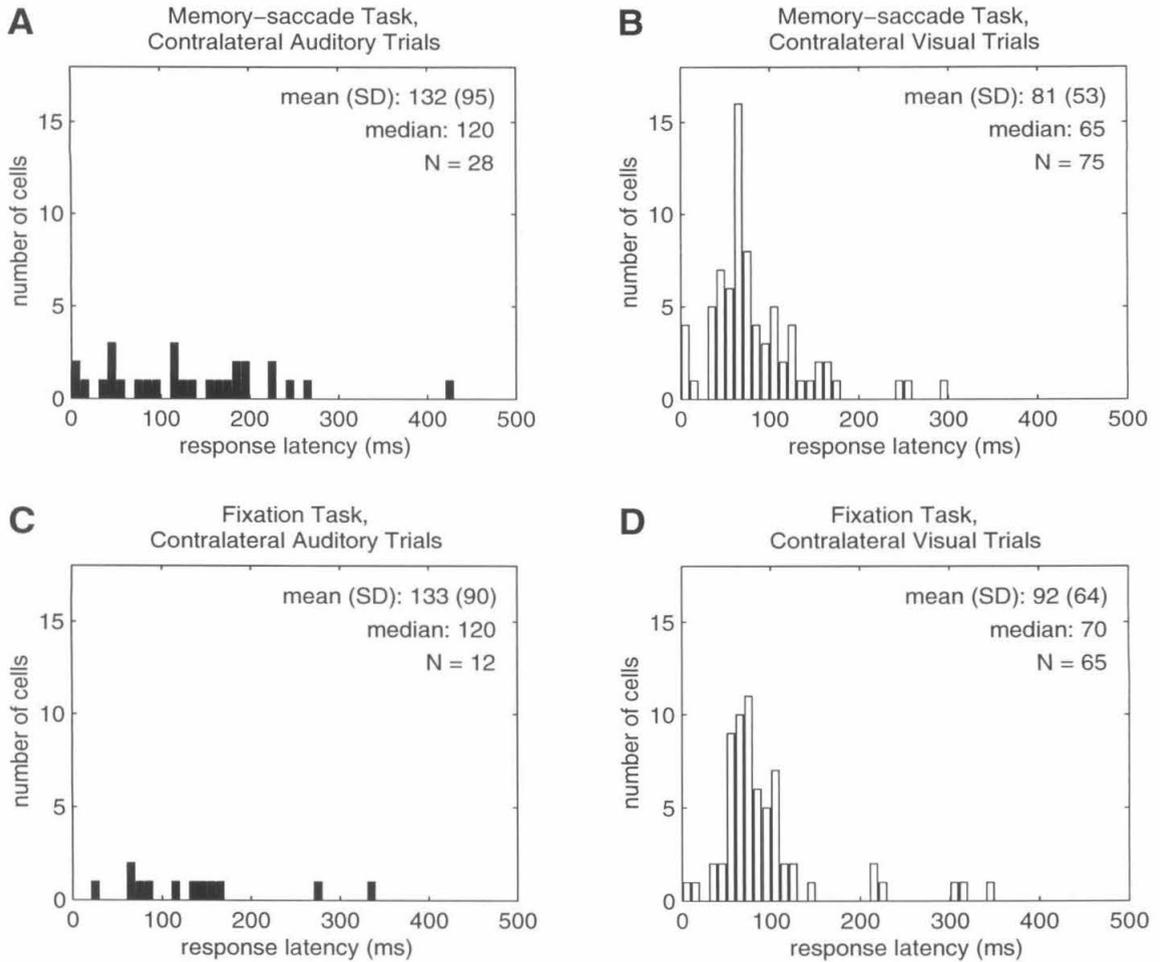


Figure 4.10: Latency histograms, for each stimulus/task condition. Response latencies were estimated for contralateral trials from all cells with spatially tuned stimulus-period activity in the indicated stimulus/task condition. Means, standard deviations, and medians of the latency distributions are shown. N is the total number of cells in each plot; histograms include only those cells with identifiable response latencies. All histogram bins are 10 ms wide.

Auditory response latencies are longer than visual response latencies not only across the population, but also in individual cells. For example, among the 20 cells with spatially tuned stimulus-period activity (and identifiable response latencies) in both the auditory and the visual memory-saccade tasks, the mean auditory response latency is 142 ± 97 ms (median 142 ms), while the mean visual response latency is 66 ± 51 ms (median 60 ms). This difference in response latency is significant (Wilcoxon signed-rank test, $p < 0.005$). A similar result is obtained in the fixation task. For the 7 cells with spatially tuned responses and identifiable response latencies in both the auditory and the visual fixation tasks, the mean auditory response latency is 144 ± 89 ms (median 130 ms), and the mean visual response latency is 65 ± 31 ms (median 75 ms). Again, this difference is significant ($p < 0.05$). Thus even for the subset of cells which respond to both auditory and visual stimuli in a given task, auditory response latencies are longer than visual response latencies.

In principle, apparent latency differences between auditory and visual responses might arise just because auditory responses tend to be weaker than visual responses. More specifically, response latency might be a function only of total stimulus-evoked activity, not of stimulus modality. To address this possibility, auditory response latencies for each task may be compared to latencies of visual responses with similar levels of evoked activity. Visual responses are included in this restricted analysis only if the difference in firing rates between the stimulus and pre-stimulus periods falls within the range observed for auditory responses in the same task. In the memory-saccade task, visual responses from 59 cells have evoked activity within the observed auditory range (-3.1 to 18.3 Hz). The mean visual response latency for these cells, 85 ± 57 ms (median 70 ms), is significantly different from the mean auditory response latency shown in Figure 4.10 *A* (Wilcoxon signed-rank test, $p < 0.05$). In the fixation task, visual response latencies in the restricted data set (mean 94 ± 57 ms; median 85 ms; $N = 35$) are not significantly different from auditory response latencies (Figure 4.10 *C*); however, both the mean and median values of the auditory response latency distribution are larger than the corresponding measures for the visual response latency distribution. Thus auditory response latencies tend to be

longer than visual response latencies even when levels of stimulus-evoked activity are similar. Latency differences between auditory and visual responses therefore cannot be attributed solely to differences in response strength.

4.3.5 Temporal features

Though an important indicator of temporal differences between auditory and visual responses in area LIP, latency measures provide no information about the overall temporal structure of responses. The algorithms outlined in the Methods section make possible a much more comprehensive analysis of temporal features in neural data. Under the assumptions of the probabilistic generative model, a collection of spike trains from G different cells can be modelled as a mixture of G randomly scaled inhomogeneous Poisson processes, each characterized by a pair of intensity profile and stability parameters. If particular response features are common to several cells, then the data might be better modelled as a mixture of fewer than G processes. As explained in the Methods section, the optimal number of mixture components M may be determined by fitting mixtures of $M = 1 \dots G$ randomly scaled inhomogeneous Poisson models to the collection of spike trains, and then selecting the simplest model which provides good fit to the data. The parameters of this optimal mixture model give a complete characterization of the data set. The intensity profiles for the model components capture the dominant temporal features in the neural responses; the stability parameters indicate the consistency of those temporal features across trials and across cells; and the mixture probabilities quantify the relative prominence of the temporal features in the data set.

Figures 4.11 and 4.12 illustrate the idea behind this approach. Figure 4.11 displays the generative mixture model for a simulated database, along with spike trains generated from the mixture model. Although this particular mixture model is intended only as a cartoon example, the parameters used to generate the simulated database were chosen to be realistic. Each plot displays one of the four intensity profiles used to generate simulated spike trains; these intensity profiles represent idealized versions

of intensity profiles observed in fits to real data from individual cells and trial conditions. The stability parameter for each model is set to 10, a value close to the average stability parameter for fits to real neural responses. Each of these four models was used to generate data sets for five simulated cells, with 10 spike trains in each data set; this number of spike trains per data set is comparable to the average number of spike trains per cell and trial condition in the actual database. Gray rasters in each plot represent simulated spike trains, and the horizontal lines separate spike trains generated from different simulated cells. The total number of simulated cells (20) is close the total number of cells with spatially tuned auditory responses in the fixation task (19), and therefore the size of the entire simulated database is similar to the size of the smallest subset of the real database which will be considered in this chapter. For the hypothetical situation in which responses recorded from 20 different cells have only four distinct temporal features, the simulated database is a realistic test case.

Figure 4.12 displays the optimal mixture model for this simulated database. The algorithm correctly determines the optimal four-component mixture model, classifying together the simulated cells generated from the same model component. The estimated intensity profiles capture the basic temporal structure of the true intensity profiles; the estimated stability parameters are close to the true stability values (see legend, Figure 4.12); and the mixture probabilities are correct. Thus the mixture-model fitting procedure successfully identifies the dominant temporal features in the simulated database.

To identify temporal features of stimulus-evoked responses in the real database, the algorithms were run separately for each of the four stimulus/task conditions, on contralateral trials recorded from neurons with spatially tuned stimulus-period responses. Thus optimal mixture models were determined for four different subsets of the entire database. These subsets consisted of contralateral trials from the 35 cells responding to auditory stimuli in the memory-saccade task, the 89 cells responding to visual stimuli in the memory-saccade task, the 19 cells responding to auditory stimuli in the fixation task, or the 73 cells responding to visual stimuli in the fixation task. As previously explained, only cells with spatially tuned responses were

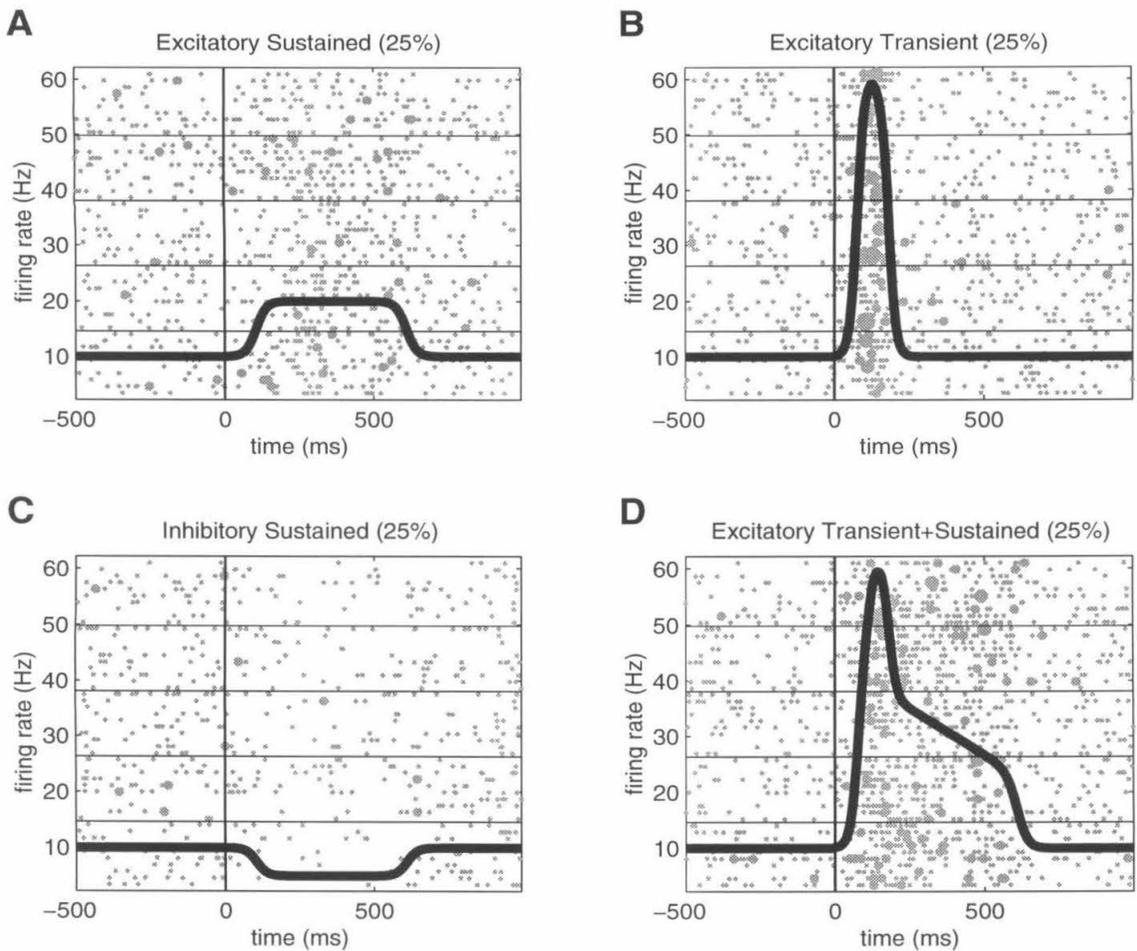


Figure 4.11: A simulated database with four temporal features, used to test algorithms for temporal feature analysis. The thick line in each plot is the intensity profile for one of the models used to generate simulated spike trains. The stability parameter is 10 for all four models. Simulated data sets for 5 cells, with 10 spike trains in each data set, were generated from each of the four models; all mixture probabilities are therefore 0.25 (25%). Gray rasters represent simulated spike trains; horizontal lines separate spike trains generated from different simulated cells. As in previous raster plots, the size of each dot is scaled to reflect the number of spikes in the corresponding 5 ms bin, and the thin vertical line indicates the time of (simulated) stimulus onset.

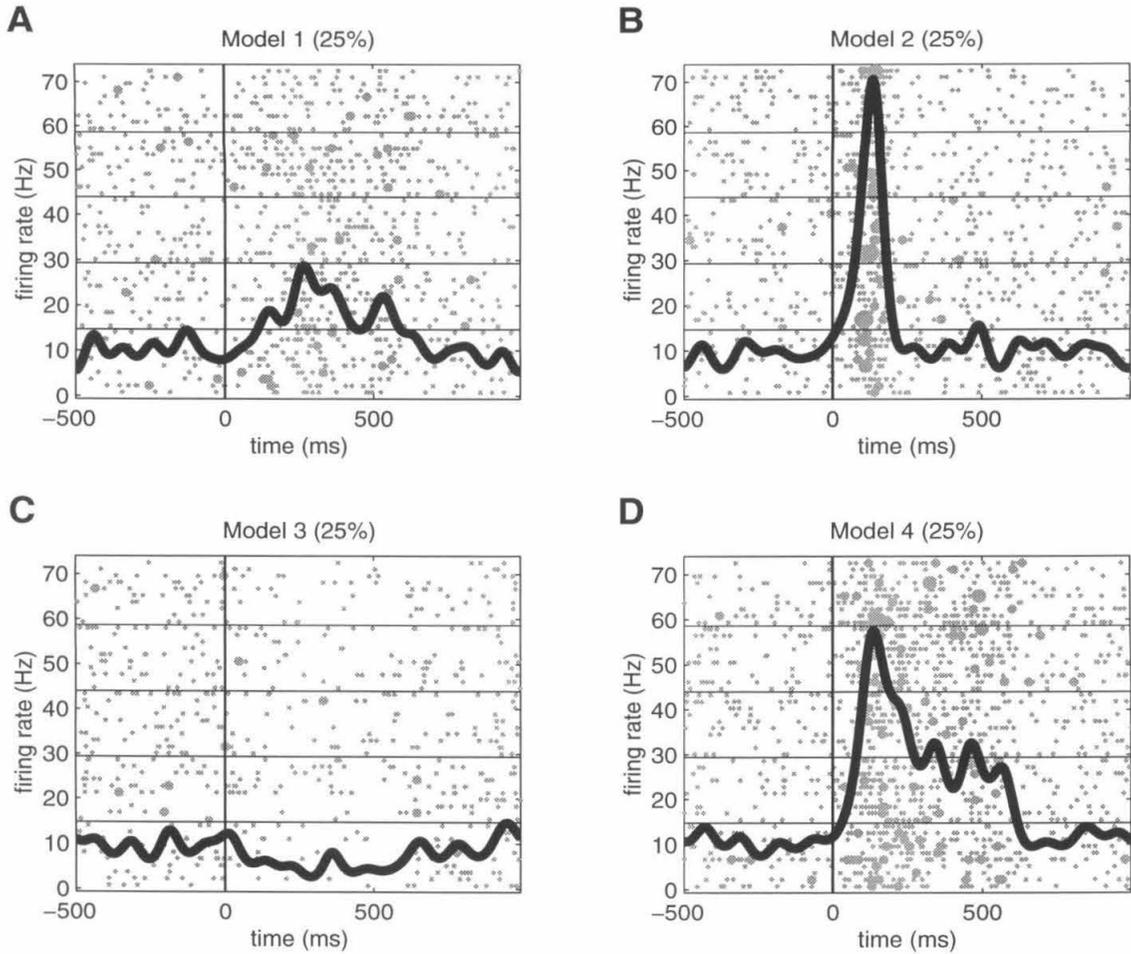


Figure 4.12: Optimal mixture model fit to the simulated database. Each panel shows the intensity profile for one of the optimized mixture components, along with the spike trains generated from simulated cells classified as belonging to that model. The percentage printed in each title indicates the mixture probability for each model. Other conventions are as in Figure 4.11. Stability parameters: *A* 16.6; *B* 15.1; *C* 12.0; *D* 12.6.

included in the analysis, because spatially untuned responses cannot be distinguished from arousal effects. The analysis database was further restricted to spike trains recorded during contralateral trials, because the vast majority of cells with spatially tuned responses showed stimulus-related modulation of their activity mainly during contralateral trials.

Confirming at the single-cell level the results observed at the population level in Figure 4.3, the temporal feature analysis demonstrates that auditory responses in the memory-saccade task tend to have gradual onsets and sustained response profiles. Figure 4.13 (pages 148 and 149) illustrates the optimal mixture model fit to the auditory memory-saccade database. Each plot shows the intensity profile for a single mixture component as a thick wavy line. Gray raster plots display the real spike trains recorded from cells assigned to each model component, and horizontal lines separate spike trains from different cells. The percentages in the titles indicate the mixture probabilities; stability parameters for the fits are listed in the figure legend.

Over one-fifth (22%) of the cells in this database have a sustained response which grows quickly from a very low baseline firing rate, and then decays through the stimulus period (Figure 4.13 *A*; page 148). An equal percentage of the cells exhibit an extremely slow response (Figure 4.13 *B*); activity builds gradually during the entire stimulus period and remains elevated until well after stimulus offset. Other neurons in this database display anticipatory activity in the pre-stimulus period, gradual response elevation after stimulus onset, peak activity in the middle of the stimulus period, and a slow decline in activation over the remainder of the stimulus period and the first 500 ms of the delay period (Figure 4.13 *C* and *D*). This temporal response profile is imposed upon a relatively high overall firing rate for approximately 13% of the cells (Figure 4.13 *C*), and on a low firing rate for another 12% of the cells (Figure 4.13 *D*). Thus a similar temporal feature emerges among cells with both high and low baseline firing rates. The remaining 31% of the cells in the auditory memory-saccade database have a very weak auditory response, riding on a relatively strong overall firing rate (Figure 4.13 *E*; page 149).

The optimal mixture model for the visual memory-saccade task (Figure 4.14;

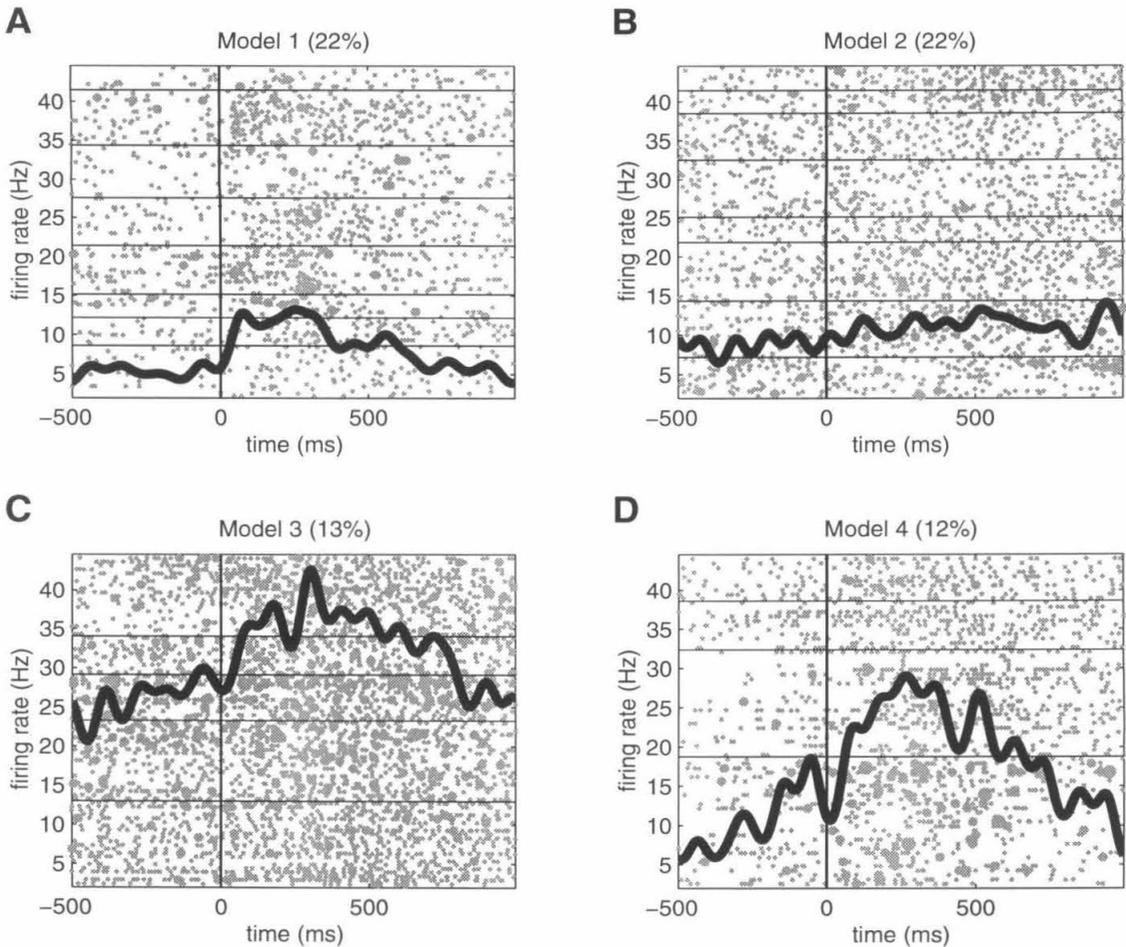


Figure 4.13: Optimal mixture model for auditory responses recorded during the memory-saccade task. Each panel displays the intensity profile for a component of the mixture model, along with spike trains of cells classified into that model component. Mixture probabilities are indicated as percentages in the titles. As in previous figures, the thin vertical line marks the time of stimulus onset, and raster dots are scaled to reflect the number of spikes in each 5 ms time bin. Horizontal lines separate spike trains recorded from different cells. Stability parameters: *A* 5.2; *B* 9.9; *C* 15.2; *D* 7.9; *E* 23.1. See page 149 for continuation of figure.

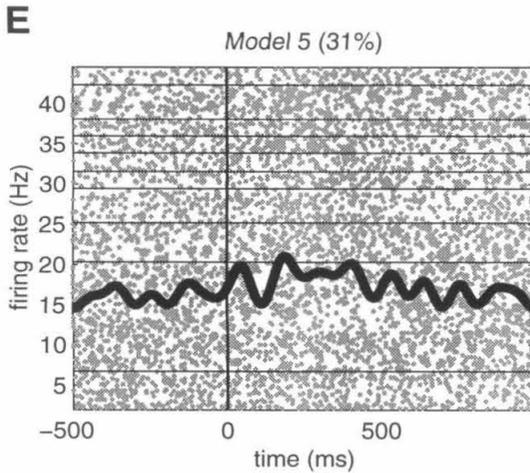


Figure 4.13: Optimal mixture model for auditory responses in the memory-saccade task, continued.

pages 150, 151, and 152) reveals stimulus-evoked responses with temporal characteristics very different from those observed in the auditory memory-saccade task. Many neurons in the visual memory-saccade database exhibit a strong transient followed by sustained activity during the stimulus period (Figure 4.14 *A–D*; page 150). This type of response profile is observed in approximately 21% of the cells total, with some variation both in the overall firing rate and in the relative strength of transient and sustained response components (for example, compare Figure 4.14 *B* and *D*). Other neurons in this database have the fast transient but no sustained response (6% of cells, Figure 4.14 *E*; page 151), or a very weak transient which leads into prolonged low-level activity (5% of cells, Figure 4.14 *F*). In addition, another 29% of the cells show weak inhibitory responses, imposed upon either a low overall firing rate (Figure 4.14 *G*) or a high firing rate (Figure 4.14 *H*). These visual response profiles are very different from the typical auditory response profiles, since temporal features extracted from the auditory memory-saccade database have neither onset transients nor inhibitory activity.

However, a small subset of the visual responses in the memory-saccade task do seem to have the same temporal characteristics as auditory responses. Approximately

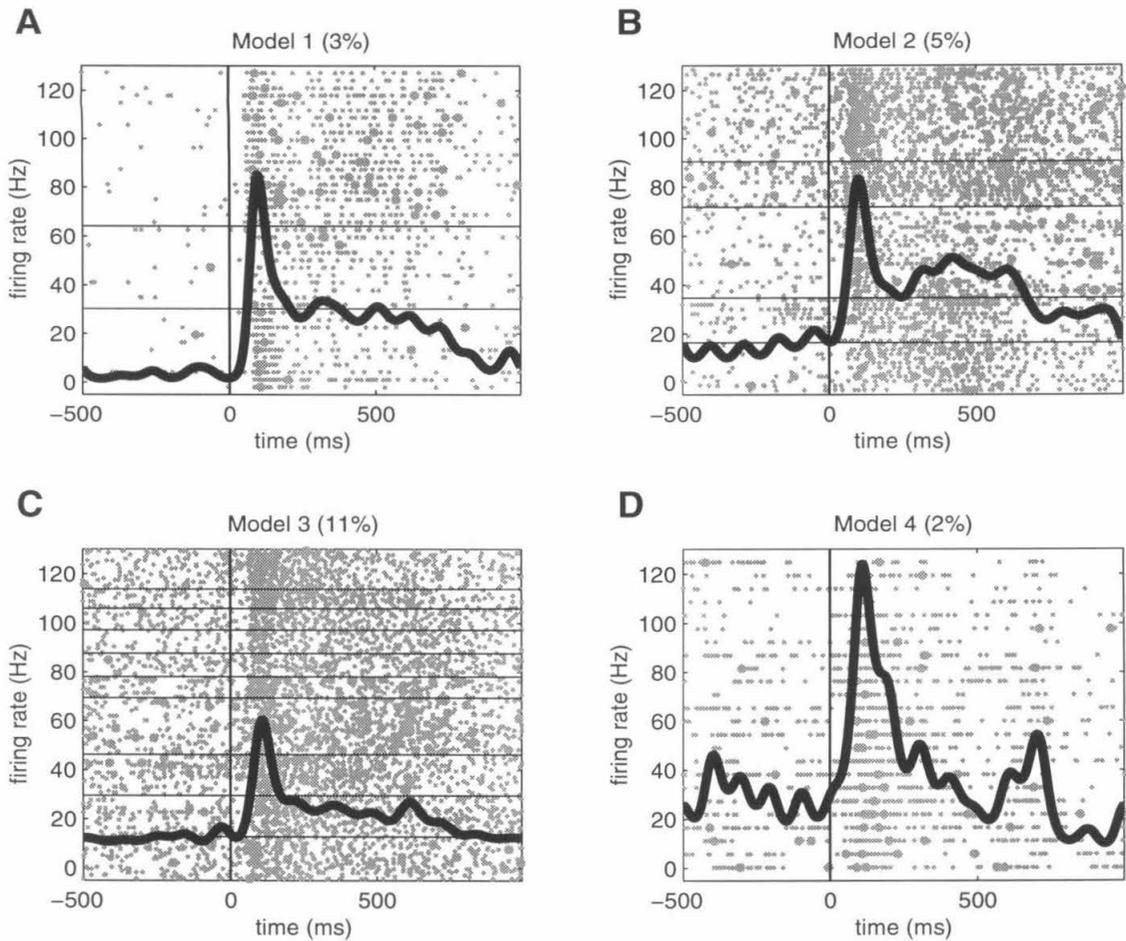


Figure 4.14: Optimal mixture model for visual responses recorded during the memory-saccade task. All conventions are as in Figure 4.13. Stability parameters: *A* 20.0; *B* 9.6; *C* 26.9; *D* 8.2; *E* 13.3; *F* 9.2; *G* 3.5; *H* 6.6; *I* 6.6; *J* 22.0; *K* 6.1; *L* 7.3. See pages 151 and 152 for continuations of figure.

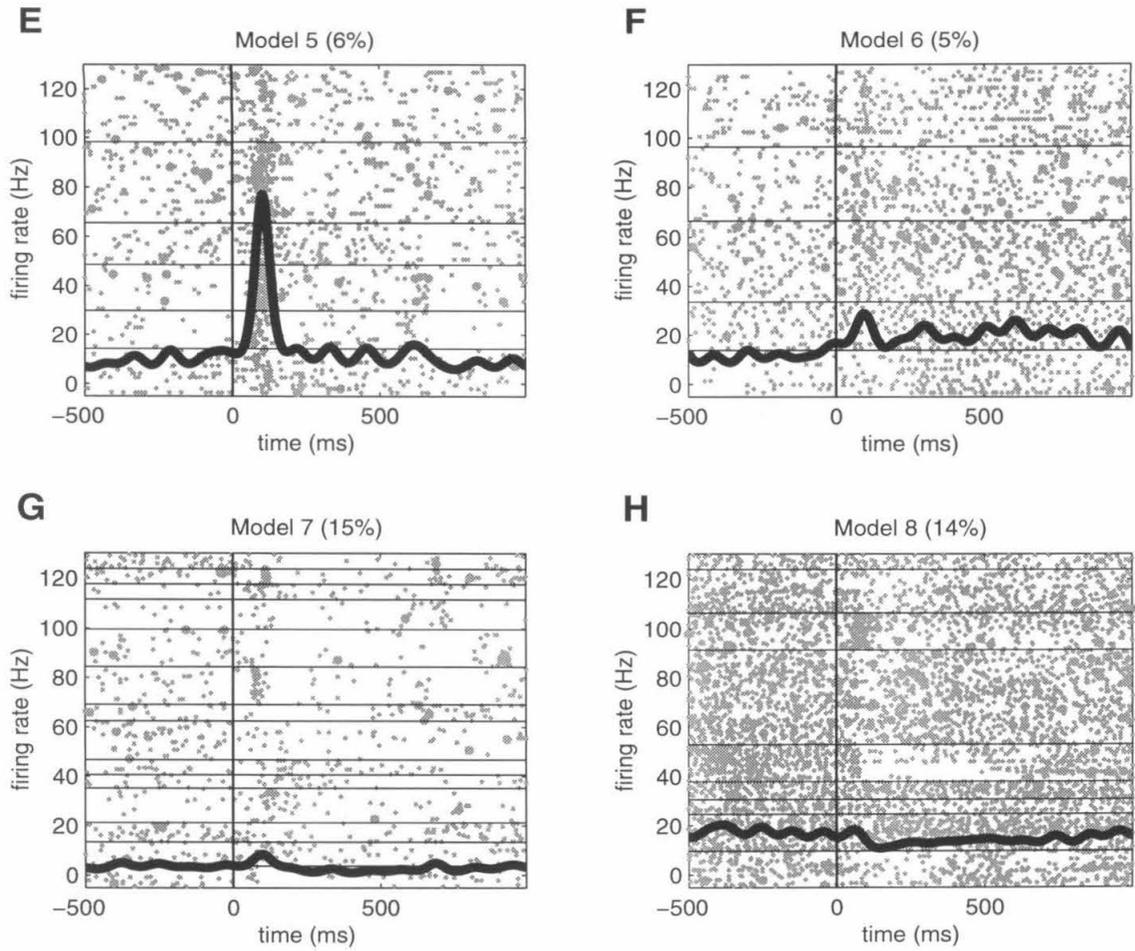


Figure 4.14: Optimal mixture model for visual responses recorded during the memory-saccade task, continued.

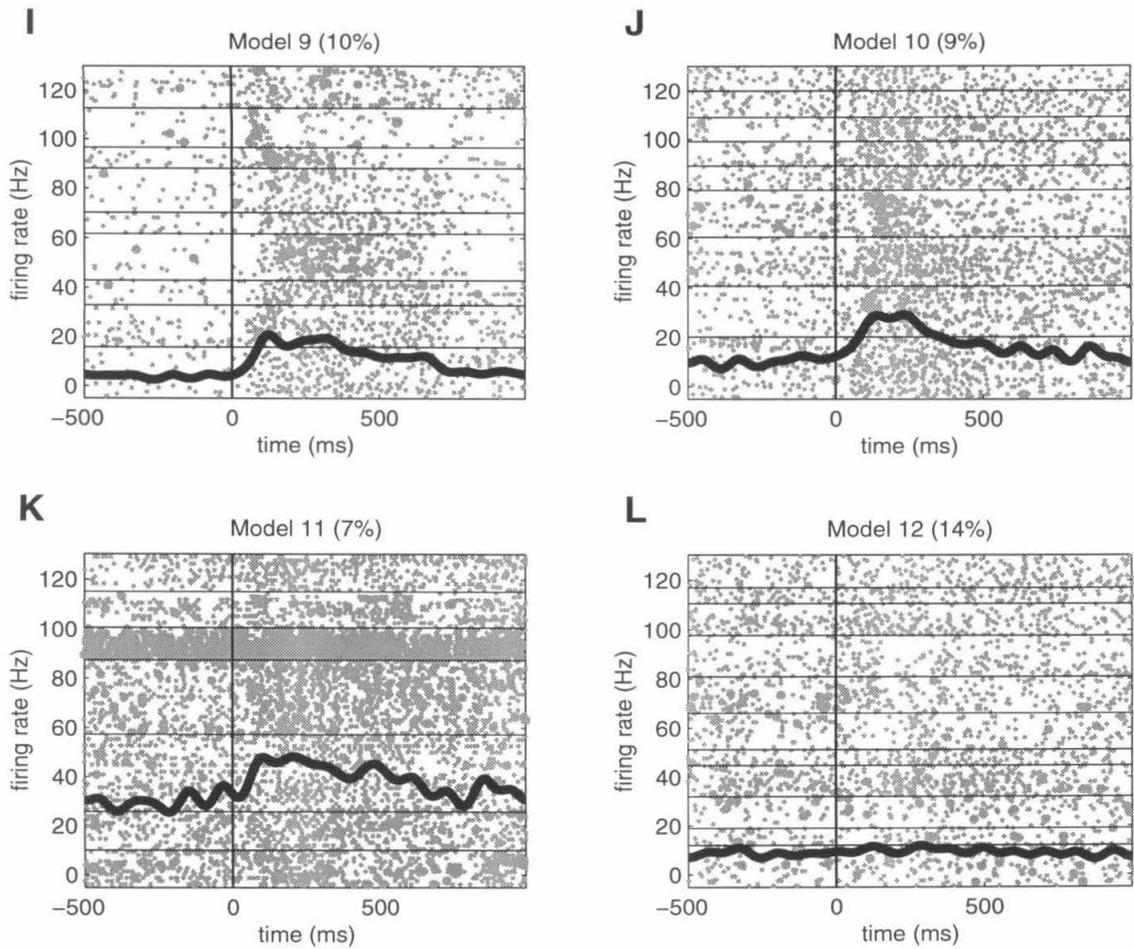


Figure 4.14: Optimal mixture model for visual responses recorded during the memory-saccade task, continued.

26% of cells in the visual memory-saccade database exhibit sustained responses which are gradual in onset (Figure 4.14 *I–K*; page 152). This basic response profile, with minor variations in the duration of sustained activity, is evident both in neurons with relatively low firing rates (10% of cells, Figure 4.14 *I*; 9% of cells, Figure 4.14 *J*), and in neurons with high baseline activity (7% of cells, Figure 4.14 *K*). The temporal features of these responses resemble those observed for auditory responses in Figure 4.13 *A*, *C*, and *D* (page 148). The remaining 14% of cells in the visual memory-saccade database (Figure 4.14 *L*) seem to have very little stimulus-evoked activity.

The differences between auditory and visual response profiles observed in the memory-saccade task are also evident in the fixation task. Figure 4.15 illustrates the optimal mixture model for the auditory fixation task. About 60% of the cells show a rise in activity after stimulus onset which quickly plateaus into a sustained response lasting for the duration of the stimulus period (Figure 4.15 *A*). Another 30% of the neurons seem to exhibit little or no stimulus-evoked activity, and are characterized by extremely low overall firing rates (Figure 4.15 *B*). The three remaining neurons (approximately 10% of the 19-cell database) have very high firing rates, and display a weak elevation in activity late in the stimulus period (Figure 4.15 *C*).

The optimal mixture model for the visual fixation task (Figure 4.16; pages 155, 156, and 157) reveals that the majority of visual responses in the fixation task have strong onset transients. Over one-third of the cells display a gradual decline in activity after this initial transient (Figure 4.16 *A*, *B*, and *C*; page 155). For a few neurons (4% of cells, Figure 4.16 *D*), the sharp transient response is followed by robust sustained activity, which remains at a relatively constant level throughout the stimulus period. Responses of many other cells exhibit only a weak transient, with little or no sustained activity (Figure 4.16 *E* and *F*; page 156). And like the visual memory-saccade database, this database contains some cells with inhibitory responses to visual stimuli (15% of cells; Figure 4.16 *G*). As noted previously, onset transients and inhibitory profiles are not observed among auditory responses.

However, like the visual memory-saccade database, this database also contains visual responses with temporal characteristics more typical of auditory responses.

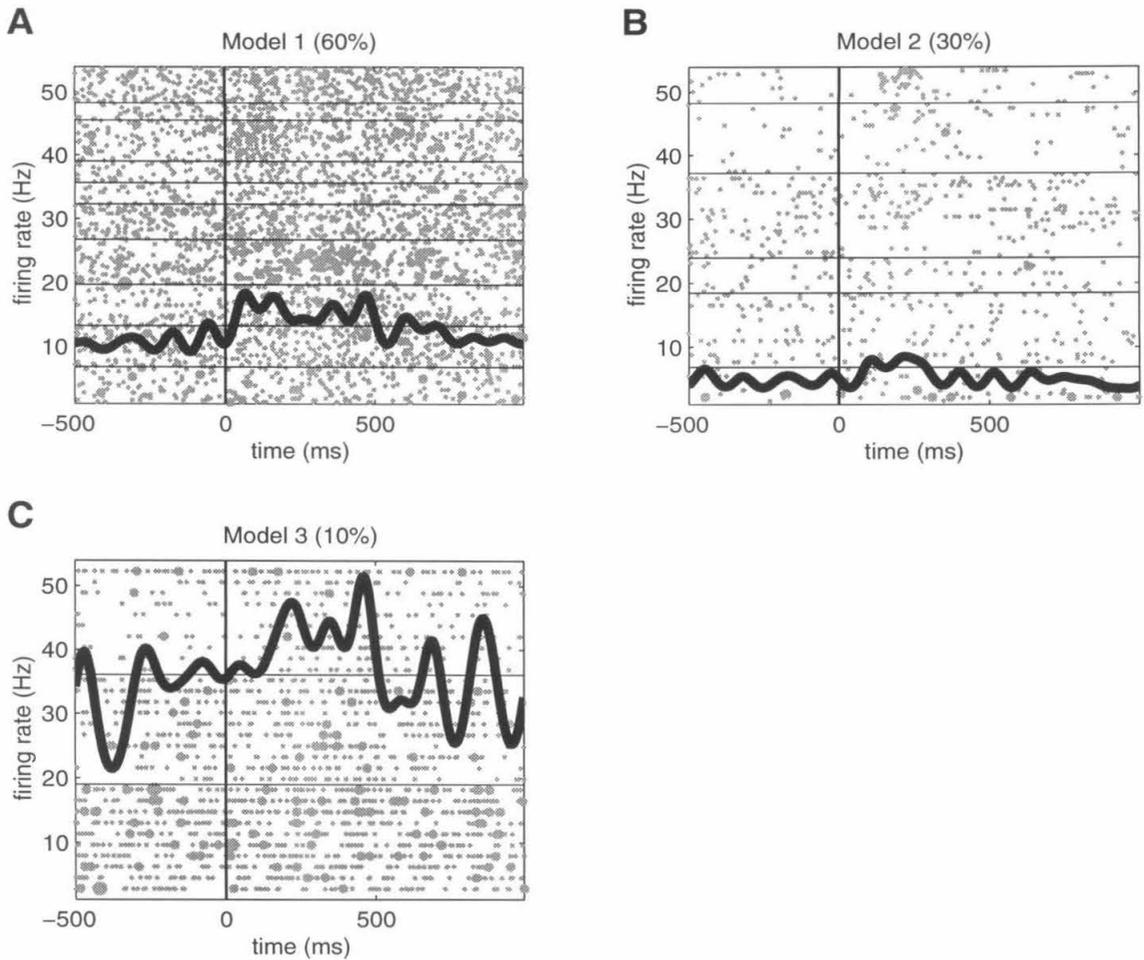


Figure 4.15: Optimal mixture model for auditory responses recorded during the fixation task. All conventions are as in Figure 4.13. Stability parameters: *A* 13.1; *B* 2.7; *C* 4.8.

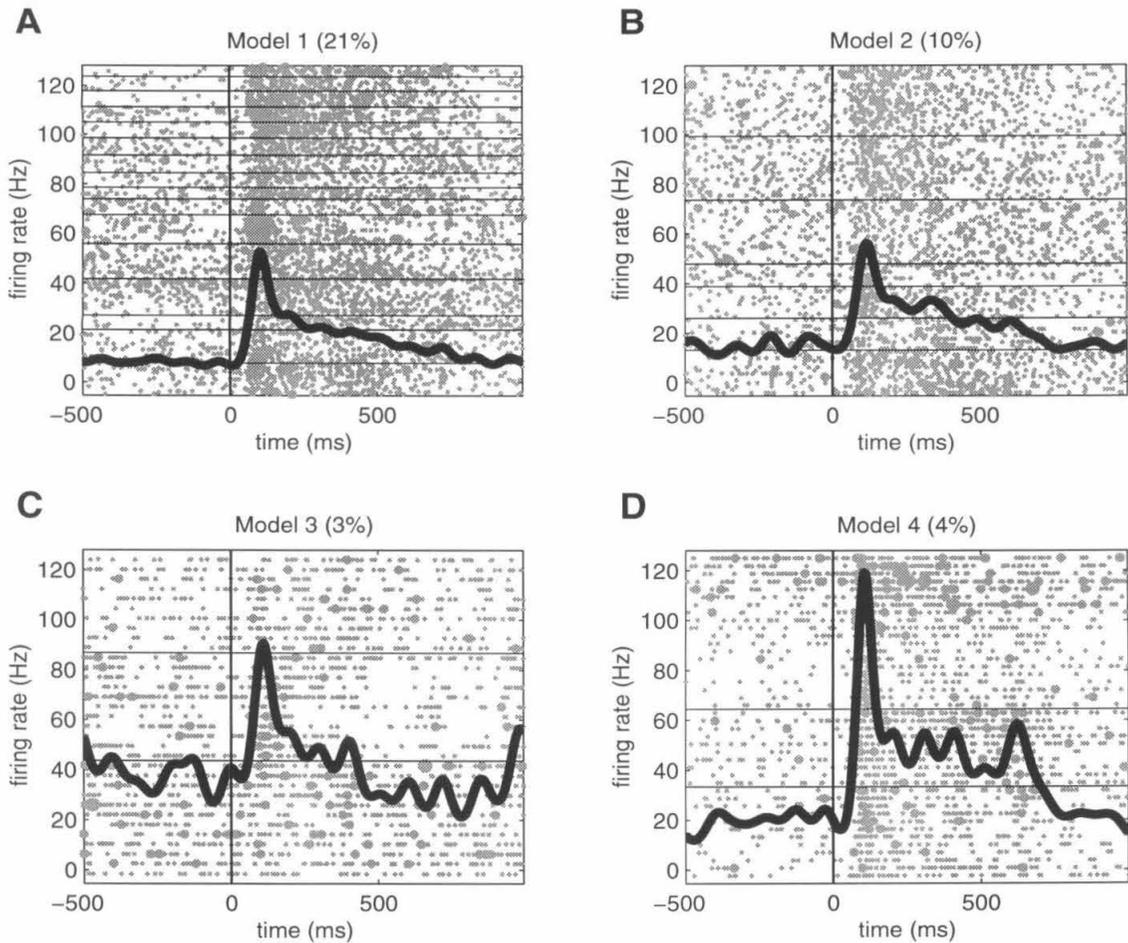


Figure 4.16: Optimal mixture model for visual responses recorded during the fixation task. All conventions are as in Figure 4.13. Stability parameters: *A* 12.0; *B* 35.0; *C* 29.3; *D* 10.0; *E* 1.6×10^4 ; *F* 2.4; *G* 5.2; *H* 29.8; *I* 8.2; *J* 38.4; *K* 1.1×10^5 . See pages 156 and 157 for continuations of figure.

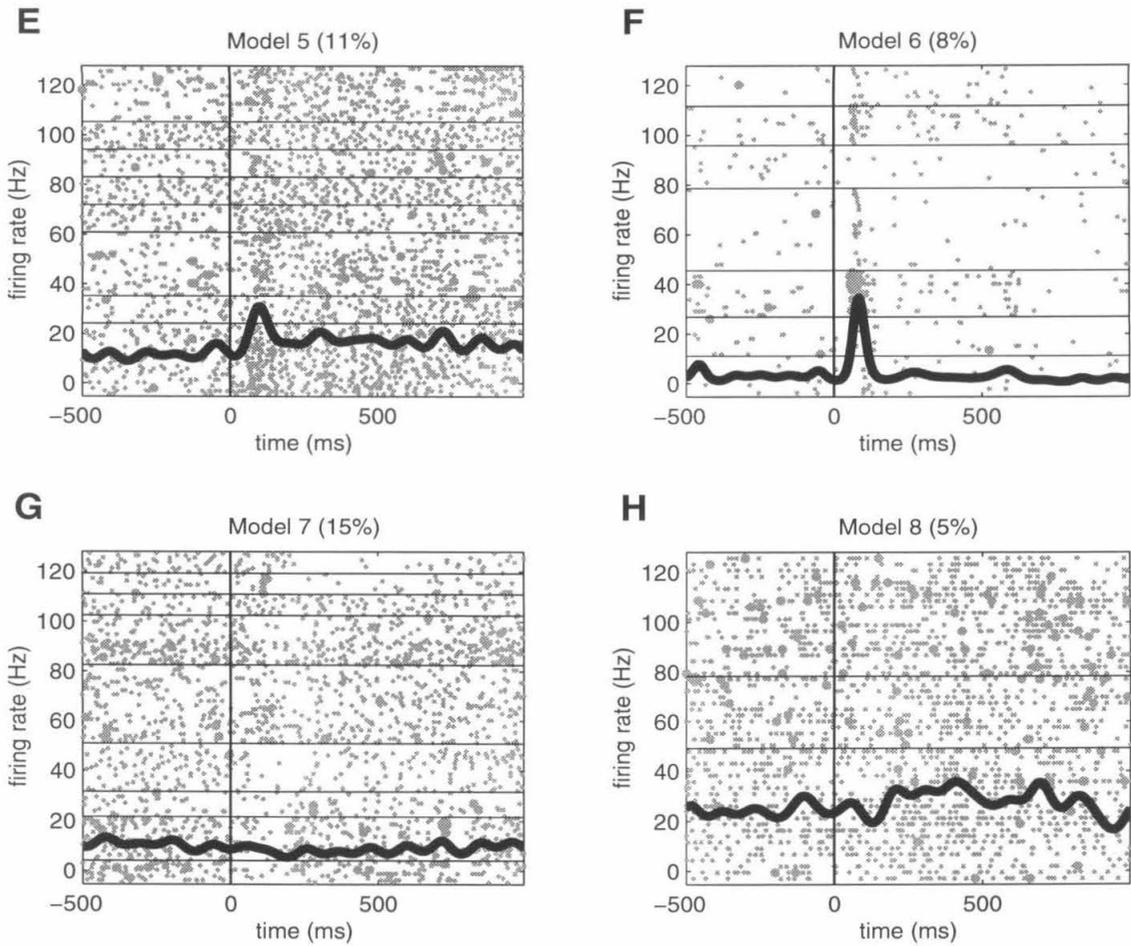


Figure 4.16: Optimal mixture model for visual responses recorded during the fixation task, continued.

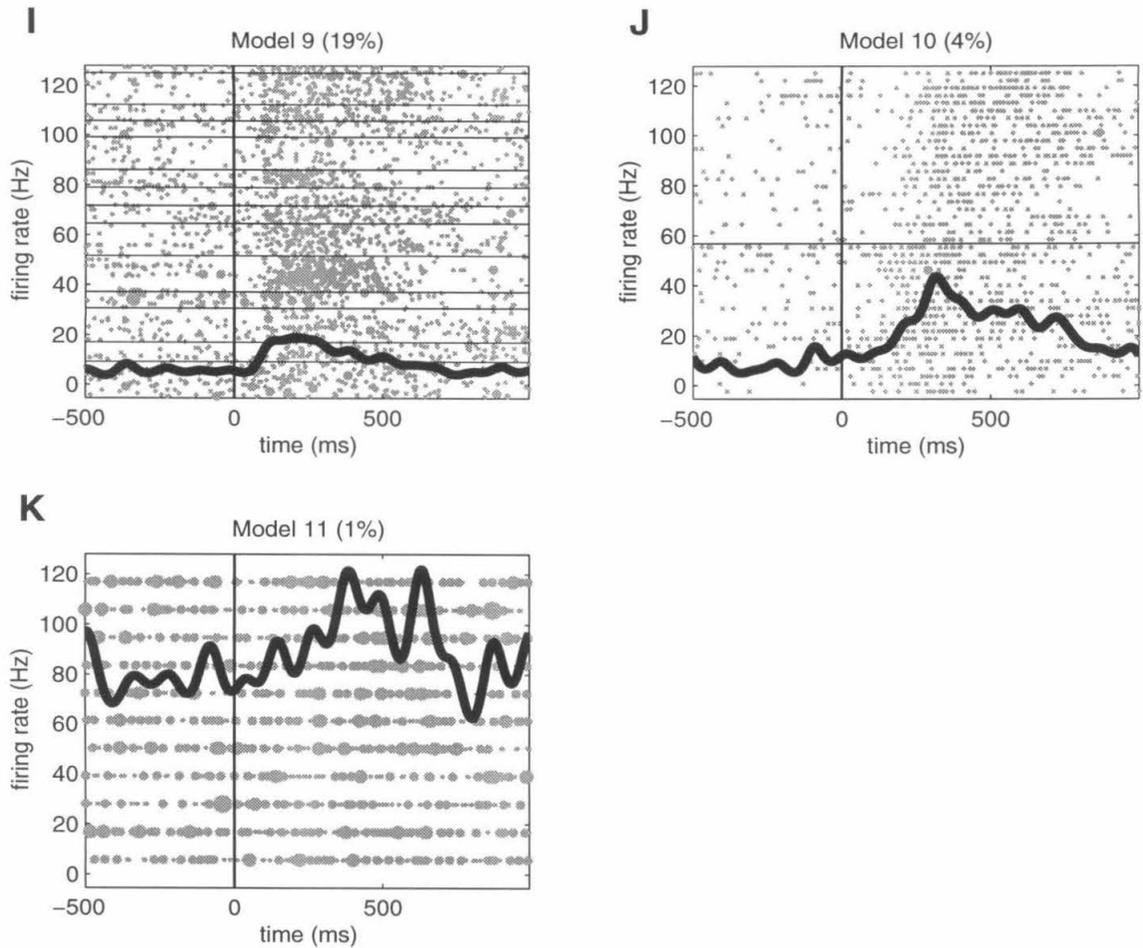


Figure 4.16: Optimal mixture model for visual responses recorded during the fixation task, continued.

About 5% of the cells display a very weak sustained elevation in activity (Figure 4.16 *H*; page 156), which resembles a temporal feature previously observed among auditory responses in Figure 4.13 *B*. Another 19% of the neurons exhibit the temporal feature perhaps most characteristic of auditory responses: a very gradual rise in activity which declines during the stimulus period (Figure 4.16 *I*; page 157). For a few other cells in the database, the visual responses are even more profoundly delayed; these neurons do not show peak activity until halfway through the stimulus period (4% of cells, Figure 4.16 *J*; 1 cell, Figure 4.16 *K*). These visual responses are similar to the auditory responses in Figure 4.15 *C*.

These four mixture models supersede the population histograms in Figure 4.3, by providing a complete and quantitative characterization of temporal features in auditory and visual responses at the single-cell level. The analysis demonstrates that auditory responses in both the memory-saccade task and the fixation task are uniformly devoid of transients, predominantly excitatory, and usually slow in onset. Visual responses, on the other hand, may be either excitatory or inhibitory; moreover, the excitatory responses usually have a fast transient component, which appears either in isolation or in combination with sustained activity. The temporal profiles most characteristic of auditory responses do sometimes appear in visual responses, but are observed only among a small subset of the visually responsive cells in each task. Thus the temporal feature analysis provides a great deal of information beyond that available in the population histograms, regarding the diversity of temporal features in neural responses at the single-cell level.

4.4 Discussion

4.4.1 Temporal features of auditory and visual responses

The analyses presented in this chapter demonstrate that the temporal features of responses to auditory stimuli in area LIP are different from the temporal features of most responses to visual stimuli. Population histograms indicate that auditory

responses are, on average, more gradual in onset than visual responses. Latency analyses reveal that auditory response latencies are significantly longer than visual response latencies, both across the population and in individual cells. Finally, and most importantly, a novel method for analyzing temporal features at the single-cell level shows that most auditory responses are weakly excitatory, with slow onsets and relatively sustained activity profiles. In contrast, visual responses can be either inhibitory or excitatory, and the excitatory responses often have a fast transient component. Visual responses which resemble auditory responses do exist, but represent a minority of all the responses to visual stimuli.

This study is the first systematic investigation of temporal features in LIP responses to auditory and visual stimuli. However, it is not the first study to examine auditory and visual response latencies in LIP. Mazzoni et al. (1996b) also analyzed LIP activity recorded during auditory and visual memory-saccade tasks, but found no significant difference between auditory and visual response latencies. Furthermore, Mazzoni et al. (1996b) reported response latencies to be about 30–50 ms longer overall than the latencies observed in the present study. Although these results appear to conflict with the findings of this chapter, three considerations suggest that the apparent discrepancies might be illusory. First of all, the median auditory response latency reported in the Mazzoni et al. (1996b) study was indeed longer than the median visual response latency (155 ms vs. 125 ms). Histograms shown in Mazzoni et al. (1996b) suggest that a similar trend exists for mean response latencies, although mean latency values were not reported in that paper. Second, Mazzoni et al. (1996b) analyzed a smaller number of cells than were considered in the present study, and therefore would have been less likely to detect small but significant differences between auditory and visual response latencies overall. Third, the method used for measuring response latencies in the Mazzoni et al. (1996b) study was quite different from the method used here; although this methodological difference would not be expected to affect comparisons between stimulus modalities within each study, it may explain the offset between the two studies in absolute latency values overall.

Assuming, then, that auditory response latencies in area LIP are indeed longer

than visual response latencies in general, the pathway by which auditory information reaches area LIP becomes an issue. In theory, the long latencies of auditory responses in area LIP might be attributed to the extra processing time required to compute sound location from interaural time and intensity differences, and monaural spectral cues. However, auditory response latencies rarely exceed 100 ms even in area Tpt (Leinonen et al. 1980), a high-level auditory association area which projects to the inferior parietal lobule (Hyvärinen 1982b) and is thought to be involved in sound localization (Kaas and Hackett 1998). Moreover, auditory response latencies are significantly shorter than visual response latencies in the deeper layers of the superior colliculus (Jay and Sparks 1987a; Wallace et al. 1996), where neurons have spatially tuned auditory as well as visual receptive fields.

Thus the slow onsets and sustained profiles of auditory responses in area LIP stand in stark contrast to the short latencies and sharp transients of most visual responses, and suggest that auditory signals enter area LIP through a much more circuitous route than visual signals. Indeed, while area LIP receives strong feedforward projections from many extrastriate visual areas (Andersen et al. 1990; Blatt et al. 1990), inputs from auditory association areas seem to be relatively sparse (Pandya and Kuypers 1969; Divac et al. 1977; Hyvärinen 1982b). However, strong anatomical projections do exist between auditory association areas and the frontal eye fields (Pandya and Kuypers 1969; Kaas and Hackett 1998; Romanski et al. 1999), and between the frontal eye fields and area LIP (Andersen et al. 1990; Blatt et al. 1990; Stanton et al. 1995). Given the distinctive temporal features of auditory responses in area LIP, it seems possible that auditory signals in LIP are not arriving via direct feedforward projections from auditory association areas at all, but rather via *feedback* projections from frontal areas involved in building associations between arbitrary stimuli and oculomotor behaviors (Chen and Wise 1995a; Chen and Wise 1995b; Chen and Wise 1996; Bichot et al. 1996).

Previous chapters of this thesis have shown that responses to auditory stimuli emerge in area LIP after auditory-saccade training, and that auditory responses are significantly more dependent on behavioral task (and more predictive of delay or

saccade activity in the memory-saccade task) than visual responses. The results of the present chapter complement the previous findings, by suggesting that responses to auditory stimuli in area LIP carry highly processed auditory information. Responses to auditory stimuli in area LIP may differ from many visual responses in the extent to which they reflect the significance of stimuli as potential saccade targets, rather than the specific sensory parameters of the stimuli. In other words, auditory responses may lie farther along the sensorimotor continuum than most visual responses in area LIP.

4.4.2 Randomly scaled inhomogeneous Poisson model

In addition to demonstrating that the temporal features of auditory and visual responses in area LIP differ, this chapter describes a novel method for spike train analysis (also discussed in Sahani 1999). This method is based on the assumption that spike trains are doubly stochastic Poisson processes. More specifically, spike trains recorded from the same cell under identical experimental conditions are assumed to arise from inhomogeneous Poisson processes, whose underlying rates are randomly scaled copies of a single smooth time-varying function. Principled algorithms for smoothing and clustering spike trains are derived from this model within the Bayesian framework, providing a solid statistical foundation for analysis of stimulus-locked temporal features in spike trains. This method therefore represents an improvement upon many existing techniques for spike train analysis, which do not provide an explicit probabilistic formulation of assumptions to guide interpretation of results (e.g., Richmond and Optican 1987; McClurkin et al. 1991; Victor and Purpura 1997; Becker and Krüger 1996; Middlebrooks et al. 1998). Moreover, unlike other explicitly probabilistic models which have previously been applied to spike train analysis (e.g., Radons et al. 1994; Abeles et al. 1995), the randomly scaled inhomogeneous Poisson model is appropriate for analysis of stimulus-locked temporal features in single-unit neural recordings.

The randomly scaled inhomogeneous Poisson model is also, to a first approxima-

tion, consistent with the statistics of spike trains recorded from LIP neurons. As is assumed in the model, LIP neurons are more variable in their firing than would be expected for simple Poisson processes, and this variability across trials usually takes the form of a multiplicative (rather than additive) scaling of an underlying Poisson rate. Indeed, goodness-of-fit tests indicate that the model achieves a good fit to the data overall, in that the likelihoods of real spike trains under the model are comparable to the likelihoods of simulated spike trains. Nevertheless, there are clear indications (see Figure 4.8) that the randomly scaled inhomogeneous Poisson model is not accurate in detail. These inaccuracies might arise at several different levels within the hierarchy of assumptions underlying the model. In particular, the following claims can be questioned.

- **Individual spike trains are instances of inhomogeneous Poisson processes.** This assumption is consistent with the conclusions of many previous studies, which have found that cortical spike trains are reasonably well modelled as inhomogeneous Poisson processes (Smith and Smith 1965; Moore et al. 1966; Tomko and Crapper 1974; Softky and Koch 1993; Shadlen and Newsome 1996; Bair and Koch 1996). However, well-known features of cortical spike trains (e.g., refractory periods and burst firing) clearly violate the Poisson assumption that interspike intervals are exponentially distributed and independent. Indeed, several investigators have identified temporal structure in cortical spike trains which could not arise from a Poisson process (Tolhurst et al. 1983; Teich et al. 1996; Baddeley et al. 1997; Victor and Purpura 1996; Victor and Purpura 1998; Reich et al. 1998). Thus the inhomogeneous Poisson assumption is certainly an oversimplification, and perhaps an extreme one. Violations of this assumption may be the major source of error in the fits of the randomly scaled inhomogeneous Poisson model to real data.
- **Inhomogeneous Poisson rate fluctuations do not exceed the temporal scale set by the intensity profile prior.** Even if the inhomogeneous Poisson assumption were correct, this next assumption of the randomly scaled

inhomogeneous Poisson model might be wrong. In the present study, the intensity profile prior acted as a low-pass filter, imposing a limit of about 50 ms on the effective temporal precision of spike trains. This temporal scale was chosen for computational reasons, and is arbitrary from a biological perspective. Indeed, the ringing apparent in some of the intensity profile estimates (e.g., Figure 4.6 *D*) suggests that the temporal resolution of the prior is too low. Evidence from other studies indicates that the stimulus-locked temporal precision of cortical spike trains *in vivo* can be very high — on order of 10 ms or less (Bair and Koch 1996). Thus the prior on the intensity profile may be unduly restrictive. (The prior on the stability parameter, while also open to question in principle, is comparatively vague and therefore unlikely to contribute greatly to errors in model fit.)

- **Temporal features in spike trains recorded from the same cell under identical experimental conditions are time-locked to external events in the trial.** This assumption is fundamental to almost all studies of stimulus-evoked activity in cortical spike trains; indeed, it is implicit in any analysis of peri-stimulus time histograms. However, it is certainly possible that spike trains analyzed in this study contain temporal structure which is not perfectly time-locked to the stimulus presentation. Response onset times, as well as overall response magnitudes, might fluctuate from trial to trial (Brody 1998). In this situation, algorithms based on the randomly scaled inhomogeneous Poisson model would fit a stimulus-locked intensity profile to the data, and then increase the estimate of trial-to-trial variability (i.e., decrease the stability parameter) to compensate for small variations in response latency from trial to trial. Thus violations of the time-locking assumption are poorly handled by the model, and may contribute to errors in model fit.
- **Trial-to-trial scaling of the intensity profile is multiplicative.** In so far as the basic Poisson assumption is correct, this claim is supported by analyses presented in the Results section. However, those analyses demonstrate only that

the multiplicative scaling hypothesis is better than one possible alternative: the additive scaling hypothesis. Moreover, the results indicate that spike trains from a small subset of cells in the database were actually better fit under the additive scaling hypothesis. The multiplicative scaling assumption is therefore an oversimplification, and a likely source for errors in the fit of the randomly scaled inhomogeneous Poisson model.

- **Scale factors for each trial are chosen at random from a gamma distribution.** This assumption was made primarily for mathematical convenience. Analyses of spike count distributions across trials (not shown) indicate that the gamma distribution is not an unreasonable choice. However, other distributions might provide a better fit to the data. Moreover, it is unlikely that trial-to-trial scaling of response profiles is purely random for real neurons; successive trials would be expected to scale more similarly than trials widely separated in time.

Given that all these assumptions of the randomly scaled inhomogeneous Poisson model are potentially flawed, it is not surprising that the fit of the model is imperfect. The advantage of the probabilistic generative model framework, however, is that these assumptions are explicit in the algorithms — and are therefore relatively easy to revise.

4.4.3 Possible improvements

Each of the possible shortcomings of the randomly scaled inhomogeneous Poisson model could be addressed in future versions of the probabilistic generative model, in order to develop improved algorithms for analysis of temporal features in spike trains. Among the possible revisions to the model:

- Each spike train might be modelled as the output of an integrate-and-fire neuron receiving balanced excitatory and inhibitory inputs, rather than as an instance of an inhomogeneous Poisson process. The intensity profile for a set of spike

trains would then be an estimate of the cell's intracellular potential as a function of time, averaged across trials. Integrate-and-fire models with balanced excitatory and inhibitory inputs can account for many of the non-Poisson characteristics of cortical spike trains (Holt et al. 1996; Reich et al. 1998; but see Softky and Koch 1993). Moreover, the relevant theoretical framework is already well-established, since integrate-and-fire models have been widely used to analyze spike train data (Fienberg 1974; Awiszus 1992; Shadlen and Newsome 1998).

- The pass-band of the prior on the intensity profile could be enlarged to allow for high temporal precision in cortical spike trains. Even better, the parameters of the intensity profile prior could be learned from the data, rather than fixed beforehand.
- The assumption that temporal features are perfectly time-locked to stimulus onset could be relaxed, to accommodate variations in response latency across trials. For example, a new parameter might be introduced, controlling time-shifts of the intensity profile from trial to trial. This time-shift parameter could be dependent on the scale factor for each trial, modelling the situation in which increased response magnitude correlates with decreased response latency. Alternatively, the time-shift parameter could be drawn independently from its own distribution for each trial, modelling the situation in which variations in response magnitude and response latency across trials are independent.
- The multiplicative scaling component of the model could be replaced with combined multiplicative and additive scaling. As explained in the Results section, there are very few cells in the database for which both the multiplicative *and* the additive scaling hypotheses could be rejected. A combined approach, in which multiplicative and additive scaling parameters are fit independently for each cell, might account for nearly all the observed variability across trials.
- Scale factors for each trial could be drawn from a Gaussian distribution trun-

cated at zero, which may provide better fit to observed distributions of spike counts than the gamma distribution (Gershon et al. 1998). In addition, it might be possible to introduce correlations between sequential draws from the scale factor distribution, to model the situation in which changes in overall cortical excitability persist across many repeated trials.

Most of these proposed improvements to the probabilistic generative model will increase the total number of parameters to be estimated. Therefore, some modifications to the mixture model fitting and selection procedures may be necessary. Even in the current implementation, the algorithms require hours of computing time to extract temporal features from any reasonably sized data set. The present mixture model fitting and selection procedure is computationally intensive because different mixture models must be fit to the data for every possible number of mixture components (see Methods); moreover, each fit must be restarted several times from different initial conditions to increase the chances of finding a global maximum in the posterior probability. The computational efficiency of the entire temporal feature analysis procedure could be improved considerably with cascading model selection techniques (Sahani 1999).

Last but not least, it should be noted that the method for analyzing temporal features described in this chapter could easily be applied to simultaneous recordings from multiple cells. In particular, the method might be useful for analyzing changes in the responses of multiple simultaneously recorded neurons over long periods of time. Given chronic multi-unit recordings obtained from the same site on different days, temporal features could be extracted separately from each data set and then compared to determine how response profiles change over time. As demonstrated in this chapter, the randomly scaled inhomogeneous Poisson method makes possible a complete and quantitative characterization of neural responses, and should therefore be applicable to a wide range of questions about temporal features in spike train data.

Chapter 5 Conclusions

The introduction posed three questions about responses to auditory stimuli in area LIP, and also raised three larger issues. The specific questions concerned the effects of training on the auditory responsiveness of area LIP; the influence of behavioral context on LIP responses to auditory stimuli; and the differences between auditory and visual responses in area LIP. Experiments described in previous chapters directly address these questions. The results of these experiments show that responses to auditory stimuli in area LIP are dependent on auditory-saccade training; that auditory responses, unlike visual responses, are stronger during a saccade task than during a fixation task; and that responses to auditory stimuli lack the fast excitatory transient frequently observed in visual responses. These findings have important implications with respect to the three larger issues raised in the introduction. First, the results suggest that area LIP plays no special role in auditory-to-oculomotor processing, in that the status of the stimulus as a saccade target seems far more critical to the involvement of area LIP than the fact that the stimulus modality is auditory. Second, the experiments indicate that while the function of area LIP may be primarily to subserve visual-to-oculomotor transformations, the area can also be recruited to assist in implementation of associations between oculomotor behaviors and non-visual stimuli. Third, the findings imply that neurophysiological studies like this one, in which monkeys made saccades to auditory stimuli in darkness with their heads immobilized, may reveal more about the capability of the brain to learn arbitrary oculomotor associations than about neural processing specific to auditory-to-oculomotor transformations. These speculations suggest several directions for future research.

5.1 Responses to auditory stimuli in area LIP

In the introduction, three questions were raised concerning responses to auditory stimuli in area LIP:

- Are responses to auditory stimuli in area LIP dependent on auditory-saccade training?
- Are auditory responses in area LIP affected by the behavioral context in which auditory stimuli appear?
- Are responses to auditory stimuli in area LIP qualitatively different from responses to visual stimuli?

Chapters 2, 3, and 4 offer answers to these questions. Results presented in those chapters demonstrate that responses to auditory stimuli in area LIP are dependent on auditory-saccade training; that auditory responses, unlike visual responses, are stronger during a saccade task than during a fixation task; and that responses to auditory stimuli lack the fast excitatory transient frequently observed in visual responses. These findings suggest that responses to auditory stimuli in area LIP are best interpreted as supramodal (cognitive or motor) responses, rather than as modality-specific sensory responses. Auditory responses in area LIP seem to reflect primarily the significance or selection of auditory stimuli as potential saccade targets. Moreover, auditory signals seem to reach LIP through a much more circuitous route than most visual signals. Auditory and visual responses in area LIP may therefore represent neural signals at different stages of sensorimotor processing.

5.2 Implications

In addition to posing the three questions listed above, the introduction also raised three larger issues:

- What is the role of area LIP in auditory-to-oculomotor processing?

- What is the function of area LIP overall?
- What is the neurophysiological basis for auditory-to-oculomotor transformation in primates?

These issues can now be considered in light of the findings of previous chapters.

5.2.1 Role of LIP in auditory-to-oculomotor processing

As explained in Chapter 1, Andersen and colleagues began their recent studies of auditory activity in area LIP with the expectation that movement-related and sensory-related response components might be partially distinguished through comparison of activity during auditory and visual memory-saccade tasks. This expectation was based on much earlier studies, which had indicated that LIP is unresponsive to passive auditory stimulation. After discovering that neurons in area LIP are active even during the stimulus period of an auditory memory-saccade task, Mazzoni et al. (1996b) proposed two possible interpretations of this activity. First, area LIP might actually be a bimodal brain area, receiving and integrating sensory information from both visual and auditory modalities. Second, area LIP might represent spatial locations and saccade targets in an abstract (supramodal) manner, independent of the sensory modality of the stimulus. Based on the observation that LIP neurons with auditory responses also tend to have visual responses (a finding replicated in the present study), Mazzoni et al. (1996b) concluded that the second scenario was most likely.

The investigations described in this dissertation corroborate this conclusion, by providing further evidence that auditory activity in area LIP is supramodal. Responses to auditory stimuli during the “sensory” stimulus period of a memory-saccade task reflect the significance of auditory stimuli as potential saccade targets. The results suggest that area LIP plays no special role in auditory-to-oculomotor processing beyond its role in oculomotor processing alone. In other words, the status of an auditory stimulus as a saccade target seems to be far more relevant to its representation in area LIP than the fact that the stimulus is a spatially localized acoustic signal. The long latencies of auditory responses further suggest that responses to auditory stimuli

in area LIP arise through feedback from frontal cortex, raising the possibility that area LIP participates in a frontal-parietal circuit subserving conditional oculomotor associations. According to this view, area LIP should be no more involved in directing eye movements to auditory stimuli than it would be in directing eye movements to any other type of stimuli associated with oculomotor behavior.

5.2.2 Function of area LIP

The ultimate implications of this reasoning are clear: area LIP might be recruited to assist in implementing associations between oculomotor behaviors and *any* stimuli. Previous studies and this dissertation demonstrate that the stimuli need not be visual, and it seems likely that the stimuli need not even be spatial. Recent reports of shape-selective LIP neurons in animals trained to make saccades to shapes (Sereno and Maunsell 1998) suggest that area LIP might represent associations between saccades and non-spatial visual cues, just as the supplementary eye fields and (to a lesser extent) the frontal eye fields do (Chen and Wise 1995a; Chen and Wise 1995b; Chen and Wise 1996). To take an extreme example, it is conceivable that training a monkey to perform rightward saccades whenever a queen-of-hearts card is presented, and leftward saccades whenever a queen-of-diamonds appears, might induce “queen-of-hearts” and “queen-of-diamonds” responses in area LIP which are indistinguishable from auditory responses.

Thus the function of area LIP might best be described as two-fold: first, to subserve visual-to-oculomotor transformations, and second, to implement associations between arbitrary stimuli and oculomotor behaviors. The first of these LIP functions is widely recognized, on the basis of anatomical and physiological evidence accumulated over many years of research (see Chapter 1). The second function is currently supported only by this dissertation and a few other studies of auditory responses in area LIP (Mazzoni et al. 1996b; Stricanne et al. 1996), which provide at best indirect evidence for the claim. Clearly, more research is necessary to determine the extent to which area LIP is involved in implementing arbitrary oculomotor associations (see

below). Future experiments aimed at exploring these speculations may provide important insight into the relationship between different aspects of LIP function. For example, area LIP might normally act as a relatively specialized visual-to-oculomotor processing module, adopting functionality in over-trained oculomotor behaviors only incidentally. Alternatively, area LIP may be involved in a much wider range of saccade behaviors, and the visual modality might turn out to have special status in area LIP only because vision provides the kind of high-precision spatial information which is most useful for guiding eye movements.

5.2.3 Auditory-to-oculomotor transformation in primates

Many previous studies (Jay and Sparks 1984; Jay and Sparks 1987b; Jay and Sparks 1987a; Vaadia et al. 1986; Schall 1991a; Schall 1991b; Russo and Bruce 1994) have investigated auditory-to-oculomotor processing under circumstances similar to those of the present experiments. More specifically, these studies have involved training monkeys to perform saccades to white-noise auditory stimuli in darkness with their heads fixed. The results of the present experiments suggest that these training circumstances may themselves influence auditory-to-oculomotor processing in primates.

Obviously, adult monkeys are capable of orienting to sounds without training under natural conditions. However, primates may not normally make eye movements to sounds in total darkness, without an accompanying head movement. The training difficulties encountered in this study certainly imply that auditory saccades made under such circumstances are very unnatural for macaque monkeys. Moreover, both the behavioral training difficulties and the neurophysiological training effects suggest that the use of such experimental conditions might affect the very process of auditory-to-oculomotor transformation under observation.

Although the auditory memory-saccade task was designed only to require orientation of the eyes to the remembered location of an auditory stimulus, the animals may have interpreted the task as requiring an arbitrary association between a vaguely lateralized acoustic signal and a very specific eye movement. Such auditory-oculomotor

associations may be no different in principle from associations between non-spatial auditory cues and eye movements. The behavioral and neurophysiological findings of the present experiments and similar studies may therefore reflect the workings of a neural circuit involved in implementing highly overtrained eye movements, not the functioning of the circuit which presumably underlies natural orienting movements toward auditory stimuli. Thus the results of this study, and many previous studies of auditory saccades in primates, may reveal more about the capabilities of the primate brain to learn arbitrary oculomotor associations, than about neural processing specific to auditory-to-oculomotor transformations.

5.3 Future directions

Many questions remain about the role of area LIP in auditory-to-oculomotor processing, the function of LIP overall, and the process of auditory-to-oculomotor transformation in primates. These questions, listed below, suggest several directions for future research.

- **To what extent are responses to auditory stimuli in area LIP associated specifically with oculomotor behavior?** As mentioned in previous chapters, the effects of training and task on responses to auditory stimuli in area LIP might reflect spatial attention, rather than oculomotor intention or oculomotor significance. In other words, auditory responses in area LIP might be associated only with the spatial location of the auditory target, not with the type of movement used to reach the target. Data obtained in the present experiments cannot be used to distinguish between these possible interpretations, because the monkeys were trained to make only one type of movement to auditory stimuli. However, previous experiments have already resolved this issue for visual stimuli, by demonstrating that neurons in area LIP respond selectively to visual stimuli which are targets for eye movements rather than arm movements (Snyder et al. 1997; Snyder et al. 1998). Similar experiments, in which animals perform either eye movements or arm movements to auditory

stimuli, should be done to confirm that auditory responses in area LIP are also associated specifically with oculomotor behavior.

- **What are the auditory capabilities of neurons in area LIP?** The experiments presented in this dissertation provide very little information about the range of auditory stimuli to which LIP neurons might be capable of responding, since the only auditory stimuli used in these experiments were high-frequency noise bursts. There may be some acoustic signals, such as macaque vocalizations, which evoke responses in area LIP even before animals have been trained to make saccades to the stimuli. Alternatively, there may be other auditory stimuli, such as tones of specific frequencies, which never elicit auditory responses from area LIP, regardless of the animal's training state or immediate behavioral task. Defining the auditory capabilities of LIP neurons could provide more insight into the possible sources of auditory input to area LIP.
- **Do neurons in area LIP respond to auditory stimuli outside of the visual field?** In this study, auditory stimuli were always presented at relatively small eccentricities. The results may therefore be valid only for auditory stimuli within the visual field (and within the oculomotor range). Given that neurons in the deep layers of the superior colliculus tend to have very peripheral auditory receptive fields (Wallace et al. 1996), and that orientation to auditory stimuli often involves a head movement (Whittington et al. 1981), it seems possible that very peripheral auditory stimuli might elicit responses in area LIP which are quite different from the auditory responses observed in the present study. In particular, responses to auditory stimuli outside of the visual field might not be associated with visual responses. Such auditory responses might also be independent of saccade training, unaffected by performance of a saccade task, and relatively abrupt in onset.
- **Do learned representations of auditory stimuli in area LIP generalize to untrained stimulus locations?** This dissertation argues that auditory responses in area LIP reflect associations between particular auditory stimuli and

specific oculomotor behaviors. While the data are certainly consistent with this interpretation, the results shown are not sufficient to prove the claim, in part because the auditory stimulus locations tested after auditory-saccade training were the same as those used during training. No attempt was made in the post-training experiments to look for responses to auditory stimuli at untrained stimulus locations. It is possible that neurons in area LIP might respond to auditory stimuli at novel locations after auditory-saccade training, and/or in the context of an auditory-saccade task. If so, auditory responses in area LIP might be better described as reflecting a general association between auditory stimuli and oculomotor behavior, not the specific oculomotor associations learned during training.

- **Are auditory responses in LIP less dependent on stimulus parameters than visual responses?** Results presented in previous chapters suggest that responses to auditory stimuli in area LIP depend primarily upon the animal's internal state, while visual responses have a much more substantial sensory component. If this conclusion is correct, then auditory responses in area LIP should be less affected by small variations in stimulus parameters than visual responses are. This possibility could be addressed directly through analysis of fluctuations in the responses of LIP neurons to pulsating auditory stimuli and flashing visual stimuli. According to the view proposed in this dissertation, auditory responses should show poor stimulus-locking, while visual responses should follow stimulus variations more closely.
- **Are the auditory and visual responses of bimodal LIP neurons modulated in the same way by eye position?** Previous studies have shown that while visual receptive fields of LIP neurons are fixed in an eye-centered reference frame, visual responses in area LIP are modulated by eye position (Andersen et al. 1987; Andersen et al. 1990). In other words, LIP neurons respond maximally to visual stimuli at a specific location relative to the point of fixation, but the magnitude of this maximal response depends on the angle of

gaze during fixation. Recently, Stricanne et al. (1996) demonstrated that most auditory responses in area LIP are also eye-centered, and additionally modulated by eye position. However, since that study was conducted in animals performing only auditory memory-saccade tasks, it is not known if auditory and visual responses in the same neuron are affected by eye position in exactly the same way. Theoretical considerations suggest that the dependencies of auditory and visual response magnitudes on eye position should be different if auditory signals have not been completely transformed from their head-centered sensory reference frame to an eye-centered reference frame before reaching area LIP (Xing et al. 1994). Results presented in this dissertation (and Mazzoni et al. 1996b) indicate that most LIP neurons which respond to auditory stimuli are also visually responsive, so it should be possible to address this issue by examining responses in single LIP neurons while animals perform auditory and visual memory-saccades from different initial fixation positions.

- **How do LIP neurons respond to combinations of auditory and visual stimuli?** In the experiments described in this dissertation, auditory and visual stimuli were always presented separately. The present results therefore provide no information about how responses to auditory and visual stimuli interact within area LIP. Previous studies have shown that neurons in the deeper layers of the primate superior colliculus respond non-linearly to combinations of auditory and visual stimuli (Wallace et al. 1996). More specifically, the collicular response depends not only on the activity which would be evoked by each component of the combination in isolation, but also on the temporal and spatial relationship between the auditory and visual stimuli. These interactions may appear as either enhancement or suppression of the predicted linear response, and have been taken as evidence that the superior colliculus is involved in multisensory integration. Given the strong anatomical and physiological links between LIP and the superior colliculus, it seems reasonable to predict that similar experiments in area LIP might produce the same results. However, the

findings of this dissertation would suggest that the influence of auditory stimuli on visual responses might be weaker in area LIP than in superior colliculus.

- **Do neurons in area LIP respond differently to auditory stimuli presented in an illuminated environment?** Psychophysical studies in humans indicate that sound localization is more accurate when performed in illuminated surroundings, even when visual features in the environment provide no clues regarding sound source location (Warren 1970; Platt and Warren 1972; Shelton and Searle 1980; Mastroianni 1982). These findings suggest that monkeys might be able to perform auditory saccades more accurately in lighted environments than in darkness. Thus the monkeys used in the present study may have had difficulty learning to perform auditory-saccade tasks in part because the training and recording sessions were conducted in a completely darkened room. Moreover, it is possible that LIP responses to auditory stimuli presented in a lighted environment might be stronger, faster, and more independent of behavioral task than the auditory responses observed in the present study.
- **Do neurons in area LIP respond differently to auditory stimuli if the monkey is allowed to move his head freely?** As in many other neurophysiological studies of saccades to auditory targets (Jay and Sparks 1984; Jay and Sparks 1987b; Jay and Sparks 1987a; Vaadia et al. 1986; Schall 1991a; Schall 1991b; Russo and Bruce 1994), the monkeys used in the present experiments performed all the behavioral tasks with their heads held steady. The training difficulties encountered in this study, and the training effects observed neurophysiologically, suggest that the monkeys adopted a more complex behavioral strategy for performing auditory saccades under these conditions than was originally anticipated. Faced with what may well be an unnatural behavioral task, the animals may simply have learned to associate vaguely localized auditory cues with very precise eye movements. If the animals had been allowed to move their heads freely, the natural auditory orienting response may have predominated over any tendency to memorize an auditory-oculomotor association.

Therefore, in unrestrained monkeys, responses to auditory stimuli might have very different properties — or indeed, might never appear in area LIP at all. Further experiments, in monkeys permitted to move their heads freely, are necessary to resolve this issue. The effects of head position on auditory responses should also be explored, if auditory responses do indeed appear in area LIP when monkeys are able to orient both head and eyes to auditory stimuli.

- **Do pinna movements affect auditory activity in area LIP?** Recent behavioral studies in cats have provided compelling evidence that pinna movements occur both in reaction to auditory stimulus presentations and in conjunction with eye movements (Populin and Yin 1998). Although equally comprehensive behavioral experiments have yet to be conducted in primates, preliminary reports suggest that the pinna movements of monkeys are similar to those of cats (Bruce et al. 1988). As explained in Chapter 3, pinna movements were not controlled in the present experiments, but are unlikely to have influenced the main results, since previous studies of LIP have found no major changes in responses to auditory stimuli when the pinnae of awake behaving monkeys are restrained (Stricanne et al. 1996). However, given that microstimulation in monkey parietal cortex can evoke pinna movements (Thier and Andersen 1998), subtle effects of pinna movement on responses to auditory stimuli in area LIP might be evident upon further investigation. For example, it is possible that pinna position exerts a gain-modulation effect on auditory responses in area LIP, just as eye position exerts a gain-modulation effect on visual responses (Andersen et al. 1987; Andersen et al. 1990).
- **Do auditory responses in area LIP resemble responses to diffuse visual stimuli?** Responses to auditory stimuli observed in this study were significantly more dependent on behavioral task, significantly more predictive of delay and saccade activity, and significantly slower in response onset than responses to visual stimuli. It is possible that these properties are not specific to auditory responses, but rather characterize all LIP responses to stimuli which are

more difficult to localize than punctate LEDs. This hypothesis could be tested through investigation of LIP responses to spatially indistinct visual stimuli. Visual stimuli created by placing diffusers over LEDs might evoke weaker, slower, and more task-dependent responses in area LIP than the LEDs alone.

- **Do neurons in area LIP respond to arbitrary stimuli associated with oculomotor behavior?** To test the hypothesis that LIP is involved in implementing arbitrary oculomotor associations, monkeys could be trained to perform directed eye movements in response to arbitrary cues presented at the fixation point. For example, high tones or red lights could be used to cue rightward saccades, and low tones or green lights could be used to cue leftward saccades. Through training, neurons in area LIP which originally responded during visually guided saccades to the right might develop selectivity for high tones or red lights at the fovea, and neurons initially tuned for leftward saccades might become responsive to low tones or green lights. In other words, neurons in area LIP, like neurons in frontal cortex (Chen and Wise 1995a; Chen and Wise 1995b; Chen and Wise 1996), might respond to arbitrary foveal stimuli in a predictive fashion after saccade training. Moreover, if the broader speculations of this chapter are correct, the characteristics of LIP responses to arbitrary foveal stimuli might resemble those of auditory responses observed in this study.
- **What regions of cortex are involved in auditory-to-oculomotor transformations?** Finally, an obvious avenue for future research is further exploration of areas in parietal, frontal, or temporal cortex which might be involved in auditory-to-oculomotor transformations. Given the results presented in this dissertation, and also the findings of Stricanne et al. (1996), it seems likely that transformation of auditory signals from sensory (head-centered) to oculomotor (eye-centered) reference frames occurs well before auditory signals reach area LIP. Studies of superior colliculus and frontal eye fields further suggest that the transformation is already largely complete by the time auditory information arrives in those areas (Jay and Sparks 1987a; Russo and Bruce 1994).

Perhaps, then, the transformation is taking place in auditory association cortex. Area Tpt seems particularly well-situated to play an important role in auditory-to-oculomotor transformations (Kaas and Hackett 1998). Previous research has shown that many neurons in area Tpt have head-centered auditory receptive fields (Leinonen et al. 1980); however, existing physiological characterizations of this region are very incomplete. Further studies of area Tpt may uncover far more complex interactions between auditory responses and eye movements than have been previously recognized. Such investigations may eventually reveal how auditory information is transformed into commands for eye movements.

Bibliography

- Abeles, M., H. Bergman, I. Gat, I. Meilijson, E. Seidemann, N. Tishby, and E. Vaadia (1995). Cortical activity flips among quasi-stationary states. *Proceedings of the National Academy of Sciences USA* 92, 8616–8620.
- Andersen, R. A. (1987). Inferior parietal lobule function in spatial perception and visuomotor integration. In F. Plum, V. B. Mountcastle, and S. R. Geiger (Eds.), *The Handbook of Physiology, Section 1: The Nervous System*, Volume V. Higher Functions of the Brain, Chapter 12, pp. 483–518. Bethesda, MD: American Physiological Society.
- Andersen, R. A. (1997). Neural mechanisms of visual motion perception in primates. *Neuron* 18, 865–872.
- Andersen, R. A., C. Asanuma, and W. M. Cowan (1985). Callosal and prefrontal associational projecting cell populations in area 7a of the macaque monkey: a study using retrogradely transported fluorescent dyes. *Journal of Comparative Neurology* 232, 443–455.
- Andersen, R. A., C. Asanuma, G. Essick, and R. M. Siegel (1990). Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *Journal of Comparative Neurology* 296, 65–113.
- Andersen, R. A., R. M. Bracewell, S. Barash, J. W. Gnadt, and L. Fogassi (1990). Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *Journal of Neuroscience* 10(4), 1176–1196.
- Andersen, R. A., G. K. Essick, and R. M. Siegel (1985). Encoding of spatial location by posterior parietal neurons. *Science* 230, 456–458.
- Andersen, R. A., G. K. Essick, and R. M. Siegel (1987). Neurons of area 7 activated by both visual stimuli and oculomotor behavior. *Experimental Brain Research* 67, 316–322.

- Andersen, R. A. and J. W. Gnadt (1989). Posterior parietal cortex. In R. Wurtz and M. Goldberg (Eds.), *The Neurobiology of Saccadic Eye Movements*, Chapter 8, pp. 315–335. Elsevier Science Publishers.
- Andersen, R. A. and V. B. Mountcastle (1983). The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex. *Journal of Neuroscience* 3(3), 532–548.
- Andersen, R. A., L. H. Snyder, D. C. Bradley, and J. Xing (1997). Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience* 20, 303–330.
- Andersen, R. A., L. H. Snyder, C.-S. Li, and B. Stricanne (1993). Coordinate transformations in the representation of spatial information. *Current Biology* 3, 171–176.
- Andersen, R. A. and D. Zipser (1988). The role of the posterior parietal cortex in coordinate transformations for visual-motor integration. *Canadian Journal of Physiology and Pharmacology* 66(4), 488–501.
- Arieli, A., A. Sterkin, A. Grinvald, and A. Aertsen (1996). Dynamics of ongoing activity: explanation of large variability in evoked cortical responses. *Science* 273, 1868–1871.
- Asanuma, C., R. A. Andersen, and W. M. Cowan (1985). The thalamic relations of the caudal inferior parietal lobule and the lateral prefrontal cortex in monkeys: divergent cortical projections from cell clusters in the medial pulvinar nucleus. *Journal of Comparative Neurology* 241, 357–381.
- Awiszus, F. (1992). Analytical reconstruction of the neuronal input current from spike train data. *Biological Cybernetics* 66, 537–542.
- Baddeley, R., L. F. Abbott, M. C. A. Booth, F. Sengpiel, T. Freeman, E. A. Wakeman, and E. T. Rolls (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society London B* 264, 1775–1783.

- Bair, W. and C. Koch (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Computation* 8, 1185–1202.
- Baizer, J. S., R. Desimone, and L. G. Ungerleider (1993). Comparison of subcortical connections of inferior temporal and posterior parietal cortex in monkeys. *Visual Neuroscience* 10, 59–72.
- Baizer, J. S., L. G. Ungerleider, and R. Desimone (1991). Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *Journal of Neuroscience* 11(1), 168–190.
- Balint, R. (1909). seelenlahmung des ‘Schauens,’ optische Ataxie, raumliche Störung der Aufmerksamkeit. *Psychiatric Neurology* 25, 51–81.
- Barash, S., R. M. Bracewell, L. Fogassi, J. W. Gnadt, and R. A. Andersen (1991a). Saccade-related activity in the lateral intraparietal area I. Temporal properties; comparison with area 7a. *Journal of Neurophysiology* 66(3), 1095–1108.
- Barash, S., R. M. Bracewell, L. Fogassi, J. W. Gnadt, and R. A. Andersen (1991b). Saccade-related activity in the lateral intraparietal area II. Spatial properties. *Journal of Neurophysiology* 66(3), 1109–1124.
- Baylis, G. C., E. T. Rolls, and C. M. Leonard (1987). Functional subdivisions of the temporal lobe neocortex. *Journal of Neuroscience* 7(2), 330–342.
- Becker, J. D. and J. Krüger (1996). Recognition of visual stimuli from multiple neuronal activity in monkey visual cortex. *Biological Cybernetics* 74, 287–298.
- Bichot, N. P., J. D. Schall, and K. G. Thompson (1996). Visual feature selectivity in frontal eye fields induced by experience in mature macaques. *Nature* 381, 697–699.
- Bisiach, E., L. Cornacchia, R. Sterzi, and G. Vallar (1984). Disorders of perceived auditory lateralization after lesions of the right hemisphere. *Brain* 107, 37–52.
- Blatt, G. J., R. A. Andersen, and G. R. Stoner (1990). Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *Journal of Comparative Neurology* 299, 421–445.

- Bracewell, R. M., P. Mazzoni, S. Barash, and R. A. Andersen (1996). Motor intention activity in the macaque's lateral intraparietal area. II. Changes of motor plan. *Journal of Neurophysiology* 76(3), 1457–1463.
- Bradley, D. C., M. Maxwell, R. A. Andersen, M. S. Banks, and K. V. Shenoy (1996). Mechanisms of heading perception in primate visual cortex. *Science* 273, 1544–1547.
- Brandt, S. A., C. J. Ploner, B.-U. Meyer, S. Leistner, and A. Villringer (1998). Effects of repetitive transcranial magnetic stimulation over dorsolateral prefrontal and posterior parietal cortex on memory-guided saccades. *Experimental Brain Research* 118, 197–204.
- Bremmer, F., A. Pouget, and K.-P. Hoffman (1998). Eye position encoding in the macaque posterior parietal cortex. *European Journal of Neuroscience* 10, 153–160.
- Brodmann, K. (1905). beitrage zur histologischen Lokalisation der Grosshirnrinde. Dritte Mitteilung: Die Rindenfelder der niederen Affen. *Journal of Psychology and Neurology* 4, 177–226.
- Brody, C. D. (1998). Slow covariations in neuronal resting potentials can lead to artefactually fast cross-correlations in their spike trains. *Journal of Neurophysiology* 80, 3345–3351.
- Brotchie, P. R., R. A. Andersen, L. H. Snyder, and S. J. Goodman (1995). Head position signals used by parietal neurons to encode locations of visual stimuli. *Nature* 375, 232–235.
- Brown, C. H., M. D. Beecher, D. B. Moody, and W. C. Stebbins (1980). Localization of noise bands by old world monkeys. *Journal of the Acoustical Society of America* 68(1), 127–132.
- Brown, C. H., T. Schessler, D. Moody, and W. Stebbins (1982). Vertical and horizontal sound localization in primates. *Journal of the Acoustical Society of America* 72(6), 1804–1811.

- Bruce, C. J. and J. A. Borden (1986). The primate frontal eye fields are necessary for predictive saccadic tracking. *Society for Neuroscience Abstracts* 12, 1086.
- Bruce, C. J., D. Burman, M. G. MacAvoy, and G. S. Russo (1988). Changes in pinna orientation accompany shifts in direction of gaze in monkeys. *Society for Neuroscience Abstracts* 14, 957.
- Brugge, J. F. (1992). An overview of central auditory processing. In A. N. Popper and R. R. Fay (Eds.), *The Mammalian Auditory Pathway: Neurophysiology*, Chapter 1, pp. 1–33. New York, NY: Springer-Verlag.
- Chafee, M. V. and P. S. Goldman-Rakic (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *Journal of Neurophysiology* 79, 2919–2940.
- Chee-Orts, M.-N. and L. M. Optican (1993). Cluster method for analysis of transmitted information in multivariate neuronal data. *Biological Cybernetics* 69, 29–35.
- Chen, L. L. and S. P. Wise (1995a). Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *Journal of Neurophysiology* 73(3), 1101–1121.
- Chen, L. L. and S. P. Wise (1995b). Supplementary eye field contrasted with the frontal eye field during acquisition of conditional oculomotor associations. *Journal of Neurophysiology* 73(3), 1122–1134.
- Chen, L. L. and S. P. Wise (1996). Evolution of directional preferences in the supplementary eye field during acquisition of conditional oculomotor associations. *Journal of Neuroscience* 16(9), 3067–3081.
- Cohen, Y. E. and E. I. Knudsen (1995). Binaural tuning of auditory units in the forebrain achistriatal gaze fields of the barn owl: local organization but no space map. *Journal of Neuroscience* 15(7), 5152–5168.
- Cohen, Y. E. and E. I. Knudsen (1998). Representation of binaural spatial cues in field L of the barn owl forebrain. *Journal of Neurophysiology* 79, 879–890.

- Cohen, Y. E., G. L. Miller, and E. I. Knudsen (1998). Forebrain pathway for auditory space processing in the barn owl. *Journal of Neurophysiology* 79, 891–902.
- Colby, C. L., J.-R. Duhamel, and M. E. Goldberg (1993). Ventral intraparietal area of the macaque: anatomic location and visual response properties. *Journal of Neurophysiology* 69(3), 902–914.
- Colby, C. L., J.-R. Duhamel, and M. E. Goldberg (1995). Oculocentric spatial representation in parietal cortex. *Cerebral Cortex* 5(5), 470–481.
- Colby, C. L., J.-R. Duhamel, and M. E. Goldberg (1996). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *Journal of Neurophysiology* 76(5), 2841–2852.
- Corbetta, M., E. Akbudak, T. E. Conturo, A. Z. Snyder, J. M. Ollinger, H. A. Drury, M. R. Linenweber, S. E. Petersen, M. E. Raichle, D. C. Van Essen, and G. L. Shulman (1998). A common network of functional areas for attention and eye movements. *Neuron* 21, 761–773.
- Cover, T. M. and J. A. Thomas (1991). *Elements of Information Theory*. New York, NY: John Wiley & Sons, Inc.
- Cox, D. R. (1955). Some statistical methods connected with series of events. *Journal of the Royal Statistical Society Series B* 17, 129–164.
- Critchley, M. (1966). *The Parietal Lobes*. New York, NY: Hafner Publishing Co.
- Cynader, M. and N. Berman (1972). Receptive-field organization of monkey superior colliculus. *Journal of Neurophysiology* 37, 187–201.
- Dean, A. F. (1981). The variability of discharge of simple cells in the cat striate cortex. *Experimental Brain Research* 44, 437–440.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society Series B* 39, 1–38.

- Deng, S.-Y., M. E. Goldberg, M. A. Segraves, L. G. Ungerleider, and M. Mishkin (1986). The effect of unilateral ablation of the frontal eye fields on saccadic performance in the monkey. In E. Keller and D. S. Zee (Eds.), *Adaptive Processes in The Visual and Oculomotor Systems*, pp. 201–208. Oxford: Pergamon Press.
- Divac, I., J. H. LaVail, P. Rakic, and K. R. Winston (1977). Heterogeneous afferents to the inferior parietal lobule of the rhesus monkey revealed by the retrograde transport method. *Brain Research* 123, 197–207.
- Duhamel, J.-R., F. Bremmer, S. BenHamed, and W. Graf (1997). Spatial invariance of visual receptive fields in parietal cortex neurons. *Nature* 389, 845–848.
- Duhamel, J.-R., C. L. Colby, and M. E. Goldberg (1998). Ventral intraparietal area of the macaque: congruent visual and somatic response properties. *Journal of Neurophysiology* 79, 126–136.
- Eskandar, E. N., L. M. Optican, and B. J. Richmond (1992). Role of inferior temporal neurons in visual memory II. Multiplying temporal waveforms related to vision and memory. *Journal of Neurophysiology* 68(4), 1296–1306.
- Eskandar, E. N., B. J. Richmond, and L. M. Optican (1992). Role of inferior temporal neurons in visual memory I. Temporal encoding of information about visual images, recalled images, and behavioral context. *Journal of Neurophysiology* 68(4), 1277–1295.
- Fienberg, S. E. (1974). Stochastic models for single neuron firing trains: a survey. *Biometrics* 30, 399–427.
- Gat, I., N. Tishby, and M. Abeles (1997). Hidden Markov modelling of simultaneously recorded cells in the associative cortex of behaving monkeys. *Network* 8(3), 297–322.
- Gawne, T. J., B. J. Richmond, and L. M. Optican (1991). Interactive effects among several stimulus parameters on the responses of striate cortical complex cells. *Journal of Neurophysiology* 66(2), 379–389.
- Gershon, E. D., M. C. Wiener, P. E. Latham, and B. J. Richmond (1998). Coding

- strategies in monkey V1 and inferior temporal cortices. *Journal of Neurophysiology* 79, 1135–1144.
- Gnadt, J. W. and R. A. Andersen (1988). Memory related motor planning activity in posterior parietal cortex of macaque. *Experimental Brain Research* 70, 216–220.
- Gnadt, J. W., R. M. Bracewell, and R. A. Andersen (1991). Sensorimotor transformation during eye movements to remembered targets. *Vision Research* 31(4), 693–715.
- Gnadt, J. W. and B. Breznen (1996). Statistical analysis of the information content in the activity of cortical neurons. *Vision Research* 36(21), 3525–3537.
- Gnadt, J. W. and L. E. Mays (1995). Neurons in monkey parietal area LIP are tuned for eye-movement parameters in three-dimensional space. *Journal of Neurophysiology* 73(1), 280–297.
- Goldberg, M. E. and M. A. Segraves (1989). The visual and frontal cortices. In R. H. Wurtz and M. E. Goldberg (Eds.), *The Neurobiology of Saccadic Eye Movements*, Chapter 7, pp. 283–313. Elsevier Science Publishers.
- Goldring, J. E., M. C. Dorris, B. D. Corneil, P. A. Ballantyne, and D. P. Munoz (1996). Combined eye-head gaze shifts to visual and auditory targets in humans. *Experimental Brain Research* 111, 68–78.
- Golomb, D., J. Hertz, S. Panzeri, A. Treves, and B. Richmond (1997). How well can we estimate the information carried in neuronal responses from limited samples? *Neural Computation* 9, 649–665.
- Goodman, S. J. and R. A. Andersen (1989). Microstimulation of a neural-network model for visually guided saccades. *Journal of Cognitive Neuroscience* 1(4), 317–326.
- Gottlieb, J. P., M. Kusunoki, and M. E. Goldberg (1998). The representation of visual salience in monkey parietal cortex. *Nature* 391, 481–484.

- Grunewald, A., J. F. Linden, and R. A. Andersen (1997). Auditory responses in LIP II: behavioral gating. *Society for Neuroscience Abstracts* 23, 16.
- Grunewald, A., J. F. Linden, and R. A. Andersen (1999). Responses to auditory stimuli in macaque lateral intraparietal area I. Effects of training. *Journal of Neurophysiology*. In press.
- Gur, M., A. Beylin, and D. M. Snodderly (1997). Response variability of neurons in primary visual cortex (V1) of alert monkeys. *Journal of Neuroscience* 17(8), 2914–2920.
- Harting, J. K., M. F. Huerta, A. J. Frankfurter, N. L. Strominger, and G. J. Royce (1980). Ascending pathways from the monkey superior colliculus: an autoradiographic analysis. *Journal of Comparative Neurology* 192, 853–882.
- Hécaen, H. and J. de Ajuriaguerra (1954). Balint's syndrome (psychic paralysis of visual fixation) and its minor forms. *Brain* 77, 373–400.
- Heilman, K. M., D. N. Pandya, E. A. Karol, and N. Geschwind (1971). Auditory inattention. *Archives of Neurology* 24, 323–325.
- Heller, J., J. A. Hertz, T. W. Kjaer, and B. J. Richmond (1995). Information flow and temporal coding in primate pattern vision. *Journal of Computational Neuroscience* 2, 175–193.
- Henry, G. H., P. O. Bishop, R. M. Tupper, and B. Dreher (1973). Orientation specificity and response variability of cells in the striate cortex. *Vision Research* 13, 1771–1779.
- Hikosaka, K., E. Iwai, H.-A. Saito, and K. Tanaka (1988). Polysensory properties of neurons in the anterior bank of the caudal superior temporal sulcus of the macaque monkey. *Journal of Neurophysiology* 60(5), 1615–1637.
- Hofman, P. M., J. G. A. Van Riswick, and A. J. Van Opstal (1998). Relearning sound localization with new ears. *Nature Neuroscience* 1(5), 417–421.
- Holt, G. R., W. R. Softky, C. Koch, and R. J. Douglas (1996). Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *Journal*

of Neurophysiology 75(5), 1806–1814.

- Hyvärinen, J. (1982a). *The Parietal Cortex of Monkey and Man*. Berlin: Springer-Verlag.
- Hyvärinen, J. (1982b). Posterior parietal lobe of the primate brain. *Physiological Reviews* 62(3), 1060–1129.
- Jay, M. F. and D. L. Sparks (1984). Auditory receptive fields in primate superior colliculus shift with changes in eye position. *Nature* 309, 345–347.
- Jay, M. F. and D. L. Sparks (1987b). Sensorimotor integration in the primate superior colliculus. I. Motor convergence. *Journal of Neurophysiology* 57(1), 22–34.
- Jay, M. F. and D. L. Sparks (1987a). Sensorimotor integration in the primate superior colliculus. II. Coordinates of auditory signals. *Journal of Neurophysiology* 57(1), 35–55.
- Jay, M. F. and D. L. Sparks (1990). Localization of auditory and visual targets for the initiation of saccadic eye movements. In M. A. Berkley and W. C. Stebbins (Eds.), *Comparative Perception*, Volume I: Basic Mechanisms, Chapter 11, pp. 351–374. John Wiley & Sons, Inc.
- Jones, B. and B. Kabanoff (1975). Eye movements in auditory space perception. *Perception & Psychophysics* 17(3), 241–245.
- Judge, S. J., B. J. Richmond, and F. C. Shu (1980). Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Research* 20, 535–537.
- Kaas, J. H. and T. A. Hackett (1998). Subdivisions of auditory cortex and levels of processing in primates. *Audiology and Neurootology* 3, 78–85.
- King, A. J. (1988). The integration of visual and auditory information in the brain. In D. M. Guthrie (Ed.), *Higher-Order Sensory Processing and Artificial Sensing*. Manchester: Manchester University Press.

- Kjaer, T. W., J. A. Hertz, and B. J. Richmond (1994). Decoding cortical neuronal signals: network models, information estimation and spatial tuning. *Journal of Computational Neuroscience* 1, 109–139.
- Knudsen, E. I., G. G. Blasdel, and M. Konishi (1979). Sound localization by the barn owl measured with the search coil technique. *Journal of Comparative Physiology* 133, 1–11.
- Knudsen, E. I. and M. S. Brainard (1995). Creating a unified representation of visual and auditory space in the brain. *Annual Review of Neuroscience* 18, 19–43.
- Knudsen, E. I., Y. E. Cohen, and T. Masino (1995). Characterization of a forebrain gaze field in the archistriatum of the barn owl: microstimulation and anatomical connections. *Journal of Neuroscience* 15(7), 5139–5151.
- Knudsen, E. I., S. Du Lac, and S. D. Esterly (1987). Computational maps in the brain. *Annual Review of Neuroscience* 10, 41–65.
- Knudsen, E. I. and P. F. Knudsen (1996). Disruption of auditory spatial working memory by inactivation of the forebrain archistriatum in barn owls. *Nature* 383, 428–431.
- Knudsen, E. I., P. F. Knudsen, and T. Masino (1993). Parallel pathways mediating both sound localization and gaze control in the forebrain and the midbrain of the barn owl. *Journal of Neuroscience* 13(7), 2837–2852.
- Koch, K. W. and J. M. Fuster (1989). Unit activity in monkey parietal cortex related to haptic perception and temporary memory. *Experimental Brain Research* 76, 292–306.
- Konishi, M. (1991). Deciphering the brain's codes. *Neural Computation* 3, 1–18.
- Konishi, M. (1992). The neural algorithm for sound localization in the owl. In *The Harvey Lectures*, 86, pp. 47–64. Wiley-Liss, Inc.
- Konishi, M. (1993). Neuroethology of sound localization in the owl. *Journal of Comparative Physiology A* 173, 3–7.

- Konishi, M., T. T. Takahashi, H. Wagner, W. E. Sullivan, and C. E. Carr (1988). Neurophysiological and anatomical substrates of sound localization in the owl. In G. M. Edelman, W. E. Gall, and W. M. Cowan (Eds.), *Auditory Function*, Chapter 24, pp. 721–745. John Wiley & Sons.
- Kurylo, D. D. (1991). Interaction of visually guided saccades with saccades induced by electrical stimulation of the posterior parietal cortex in the monkey. *Vision Research* 31(12), 2065–2073.
- Kurylo, D. D. and A. A. Skavenski (1991). Eye movements elicited by electrical stimulation of area PG in the monkey. *Journal of Neurophysiology* 65(6), 1243–1253.
- Kustov, A. A. and D. L. Robinson (1996). Shared neural control of attentional shifts and eye movements. *Nature* 384, 74–77.
- Lackner, J. R. (1973). Visual rearrangement affects auditory localization. *Neuropsychologia* 11, 29–32.
- Leinonen, L., J. Hyvärinen, and A. R. A. Sovijärvi (1980). Functional properties of neurons in the temporo-parietal association cortex of awake monkey. *Experimental Brain Research* 39, 203–215.
- Lewald, J. (1998). The effect of gaze eccentricity on perceived sound direction and its relation to visual localization. *Hearing Research* 115, 206–216.
- Lewald, J. and W. H. Ehrenstein (1996a). Auditory-visual shift in localization depending on gaze direction. *NeuroReport* 7, 1929–1932.
- Lewald, J. and W. H. Ehrenstein (1996b). The effect of eye position on auditory lateralization. *Experimental Brain Research* 108, 473–485.
- Li, C.-S. (1996). *Macaque Lateral Intraparietal Area and Oculomotor Behaviors*. Ph. D. thesis, California Institute of Technology, Pasadena, CA.
- Li, C.-S. R., P. Mazzoni, and R. A. Andersen (1999). Effect of reversible inactivation of macaque lateral intraparietal area on visual and memory saccades. *Journal of Neurophysiology* 81, 1827–1838.

- Linden, J. F., A. Grunewald, and R. A. Andersen (1996). Auditory sensory responses in area LIP? *Society for Neuroscience Abstracts* 22, 1062.
- Linden, J. F., A. Grunewald, and R. A. Andersen (1997). Auditory responses in LIP I: training effects. *Society for Neuroscience Abstracts* 23, 16.
- Linden, J. F., A. Grunewald, and R. A. Andersen (1998). Supramodal stimulus responses in macaque area LIP. *Society for Neuroscience Abstracts* 24, 521.
- Linden, J. F., A. Grunewald, and R. A. Andersen (1999). Responses to auditory stimuli in macaque lateral intraparietal area II. Behavioral modulation. *Journal of Neurophysiology*. In press.
- Lynch, J. C. (1992). Saccade initiation and latency deficits after combined lesions of the frontal and posterior eye fields in monkeys. *Journal of Neurophysiology* 68(5), 1913–1916.
- Lynch, J. C., A. M. Graybiel, and L. J. Lobeck (1985). The differential projection of two cytoarchitectonic subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus. *Journal of Comparative Neurology* 235, 241–254.
- Lynch, J. C. and J. W. McLaren (1989). Deficits of visual attention and saccadic eye movements after lesions of parietooccipital cortex in monkeys. *Journal of Neurophysiology* 61(1), 74–90.
- Lynch, J. C., V. B. Mountcastle, W. H. Talbot, and T. C. T. Yin (1977). Parietal lobe mechanisms for directed visual attention. *Journal of Neurophysiology* 40(2), 362–389.
- Mastroianni, G. R. (1982). The influence of eye movements and illumination on auditory localization. *Perception & Psychophysics* 31(6), 581–584.
- Maunsell, J. H. R. and D. C. Van Essen (1983). The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *Journal of Neuroscience* 3, 2563–2586.

- May, J. G. and R. A. Andersen (1986). Different patterns of corticopontine projections from separate cortical fields within the inferior parietal lobule and dorsal prelunate gyrus of the macaque. *Experimental Brain Research* 63, 265–278.
- Mazzoni, P. (1994). *Spatial Perception and Movement Planning in The Posterior Parietal Cortex*. Ph. D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Mazzoni, P., R. M. Bracewell, S. Barash, and R. A. Andersen (1996a). Motor intention activity in the macaque's lateral intraparietal area I. Dissociation of motor plan from sensory memory. *Journal of Neurophysiology* 76(3), 1439–1455.
- Mazzoni, P., R. M. Bracewell, S. Barash, and R. A. Andersen (1996b). Spatially tuned auditory responses in area LIP of macaques performing delayed memory saccades to acoustic targets. *Journal of Neurophysiology* 75(3), 1233–1241.
- McClurkin, J. W., T. J. Gawne, L. M. Optican, and B. J. Richmond (1991). Lateral geniculate neurons in behaving primates. II. Encoding of visual information in the temporal shape of the response. *Journal of Neurophysiology* 66(3), 794–808.
- McClurkin, J. W. and L. M. Optican (1996). Primate striate and prestriate cortical neurons during discrimination I. Simultaneous temporal encoding of information about color and pattern. *Journal of Neurophysiology* 75(1), 481–495.
- McClurkin, J. W., L. M. Optican, B. J. Richmond, and T. J. Gawne (1991). Concurrent processing and complexity of temporally encoded neuronal messages in visual perception. *Science* 253, 675–677.
- McClurkin, J. W., J. A. Zarbock, and L. M. Optican (1996). Primate striate and prestriate cortical neurons during discrimination II. Separable temporal codes for color and pattern. *Journal of Neurophysiology* 75(1), 496–507.
- Mechler, F., J. D. Victor, K. P. Purpura, and R. Shapley (1998). Robust temporal coding of contrast by V1 neurons for transient but not for steady-state stimuli. *Journal of Neuroscience* 18(16), 6583–6598.

- Middlebrooks, J. C., A. E. Clock, L. Xu, and D. M. Green (1994). A panoramic code for sound location by cortical neurons. *Science* 164, 842–844.
- Middlebrooks, J. C. and D. M. Green (1991). Sound localization by human listeners. *Annual Review of Psychology* 42, 135–159.
- Middlebrooks, J. C., L. Xu, C. Eddins, and D. M. Green (1998). Codes for sound-source location in nontopographic auditory cortex. *Journal of Neurophysiology* 80, 863–881.
- Moore, G. P., D. H. Perkel, and J. P. Segundo (1966). Statistical analysis and functional interpretation of neuronal spike data. *Annual Review of Physiology* 28, 496–522.
- Mountcastle, V. B., J. C. Lynch, A. Georgeopoulos, H. Sakata, and C. Acuna (1975). Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *Journal of Neurophysiology* 38(2), 871–908.
- Müri, R. M., A.-I. Vermersch, S. Rivaud, B. Gaymard, and C. Pierrot-Deseilligny (1996). Effects of single-pulse transcranial magnetic stimulation over the prefrontal and posterior parietal cortices during memory-guided saccades in humans. *Journal of Neurophysiology* 76(3), 2102–2106.
- Mushiake, H., N. Fuji, and J. Tanji (1999). Microstimulation of the lateral wall of the intraparietal sulcus compared with the frontal eye field during oculomotor tasks. *Journal of Neurophysiology* 81, 1443–1448.
- Optican, L. M., T. J. Gawne, B. J. Richmond, and P. J. Joseph (1991). Unbiased measures of transmitted information and channel capacity from multivariate neuronal data. *Biological Cybernetics* 65, 305–310.
- Optican, L. M. and B. J. Richmond (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *Journal of Neurophysiology* 57(1), 162–178.
- Oram, M. W., R. Lestienne, M. C. Wiener, and B. J. Richmond (1998). Accurately

- predicting precisely replicating spike patterns in neural responses of monkey striate cortex and LGN. *Society for Neuroscience Abstracts* 24, 1258.
- Pandya, D. N. and H. G. J. M. Kuypers (1969). Cortico-cortical connections of the rhesus monkey. *Brain Research* 13, 13–36.
- Pandya, D. N. and E. H. Yeterian (1985). Architecture and connections of cortical association areas. In A. Peters and E. G. Jones (Eds.), *Cerebral Cortex*, Volume 4: Association and auditory cortices, Chapter 1, pp. 3–61. New York, NY: Plenum Press.
- Perkel, D. H., G. L. Gerstein, and G. P. Moore (1967). Neuronal spike trains and stochastic point processes I. The single spike train. *Biophysical Journal* 7, 391–418.
- Pinek, B., J.-R. Duhamel, C. Cavé, and M. Brouchon (1989). Audio-spatial deficits in humans: differential effects associated with left versus right hemisphere parietal damage. *Cortex* 25, 175–186.
- Platt, B. B. and D. H. Warren (1972). Auditory localization: the importance of eye movements and a textured visual environment. *Perception & Psychophysics* 12(2B), 245–248.
- Platt, M. L. and P. W. Glimcher (1997a). Representation of prior probability and movement metrics by area LIP neurons. *Society for Neuroscience Abstracts* 23, 16.
- Platt, M. L. and P. W. Glimcher (1997b). Responses of intraparietal neurons to saccadic targets and visual distractors. *Journal of Neurophysiology* 78(3), 1574–1589.
- Platt, M. L. and P. W. Glimcher (1998). Response fields of intraparietal neurons quantified with multiple saccadic targets. *Experimental Brain Research* 121, 65–75.
- Populin, L. C. and T. C. T. Yin (1998). Pinna movements of the cat during sound localization. *Journal of Neuroscience* 18(11), 4233–4243.

- Quintana, J. and J. M. Fuster (1993). Spatial and temporal factors in the role of prefrontal and parietal cortex in visuomotor integration. *Cerebral Cortex* 3, 122–132.
- Radons, G., J. D. Becker, B. Dülfer, and J. Krüger (1994). Analysis, classification, and coding of multielectrode spike trains with hidden markov models. *Biological Cybernetics* 71, 359–373.
- Rauschecker, J. P., B. Tian, T. Pons, and M. Mishkin (1997). Serial and parallel processing in rhesus monkey auditory cortex. *Journal of Comparative Neurology* 382, 89–103.
- Recanzone, G. H., M. M. Merzenich, and W. M. Jenkins (1992). Frequency discrimination training engaging a restricted skin surface results in an emergence of a cutaneous response zone in cortical area 3a. *Journal of Neurophysiology* 67(5), 1057–1070.
- Reich, D. S., J. D. Victor, and B. W. Knight (1998). The power ratio and the interval map: spiking models and extracellular recordings. *Journal of Neuroscience* 18(23), 10090–10104.
- Richmond, B. J. and L. M. Optican (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. II. Quantification of response waveform. *Journal of Neurophysiology* 57(1), 147–161.
- Richmond, B. J. and L. M. Optican (1990). Temporal encoding of two-dimensional patterns by single units in primary visual cortex II. Information transmission. *Journal of Neurophysiology* 64(2), 370–380.
- Richmond, B. J., L. M. Optican, M. Podell, and H. Spitzer (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. I. Response characteristics. *Journal of Neurophysiology* 57(1), 132–146.
- Richmond, B. J., L. M. Optican, and H. Spitzer (1990). Temporal encoding of two-dimensional patterns by single units in primate primate visual cortex. I. Stimulus-response relations. *Journal of Neurophysiology* 64(2), 351–369.

- Rizzolatti, G., L. Riggio, and B. M. Sheliga (1994). Space and selective attention. In C. Umiltà and M. Moscovitch (Eds.), *Attention and Performance*, Volume XV, Chapter 9, pp. 231–265. Cambridge, MA: MIT Press.
- Robinson, D. L., M. E. Goldberg, and G. B. Stanton (1978). Parietal association cortex in the primate: sensory mechanisms and behavioral modulations. *Journal of Neurophysiology* 41(4), 910–932.
- Robinson, D. L. and J. W. McClurkin (1989). The visual superior colliculus and pulvinar. In R. H. Wurtz and M. E. Goldberg (Eds.), *The Neurobiology of Saccadic Eye Movements*, Chapter 9, pp. 337–360. Elsevier Science Publishers.
- Romanski, L. M., J. F. Bates, and P. S. Goldman-Rakic (1999). Auditory belt and parabelt projections to the prefrontal cortex in the rhesus monkey. *Journal of Comparative Neurology* 403, 141–157.
- Rose, D. (1979). An analysis of the variability of unit activity in the cat's visual cortex. *Experimental Brain Research* 37, 595–604.
- Russo, G. S. and C. J. Bruce (1994). Frontal eye field activity preceding aurally guided saccades. *Journal of Neurophysiology* 71(3), 1250–1253.
- Sahani, M. (1999). *Latent Variable Models for Neural Data Analysis*. Ph. D. thesis, California Institute of Technology, Pasadena, CA.
- Sakata, H., M. Taira, A. Murata, and S. Mine (1995). Neural mechanisms of visual guidance of hand action in the parietal cortex of the monkey. *Cerebral Cortex* 5(5), 430–447.
- Salinas, E. and L. F. Abbott (1996). A model of multiplicative neural responses in parietal cortex. *Proceedings of the National Academy of Sciences USA* 93, 11956–11961.
- Schaafsma, S. J. and J. Duysens (1996). Neurons in the ventral intraparietal area of awake macaque monkey closely resemble neurons in the dorsal part of the medial superior temporal area in their responses to optic flow patterns. *Journal of Neurophysiology* 76(6), 4056–4068.

- Schall, J. D. (1991a). Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: comparison with supplementary eye fields. *Journal of Neurophysiology* 66(2), 559–579.
- Schall, J. D. (1991b). Neuronal activity related to visually guided saccadic eye movements in the supplementary motor area of rhesus monkeys. *Journal of Neurophysiology* 66(2), 530–558.
- Schiller, P. H. (1977). The effect of superior colliculus ablation on saccades elicited by cortical stimulation. *Brain Research* 122, 154–156.
- Schiller, P. H. and I. Chou (1998). The effects of frontal eye field and dorsomedial frontal cortex lesions on visually guided eye movements. *Nature Neuroscience* 1(3), 248–253.
- Schiller, P. H., J. H. Sandell, and J. H. R. Maunsell (1987). The effect of frontal eye field and superior colliculus lesions on saccadic latencies. *Journal of Neurophysiology* 57(4), 1033–1049.
- Schiller, P. H., S. D. True, and J. L. Conway (1980). Deficits in eye movements following frontal eye field and superior colliculus ablations. *Journal of Neurophysiology* 44, 1175–1189.
- Schmitt, J. (1983). The design and construction of anechoic chambers. Master's thesis, University of Texas at Austin.
- Schwartz, G. (1978). Estimating the dimension of a model. *Annals of Statistics* 6, 461–464.
- Seidemann, E., I. Meilijson, M. Abeles, H. Bergman, and E. Vaadia (1996). Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *Journal of Neuroscience* 16(2), 752–768.
- Seltzer, B. and D. N. Pandya (1984). Further observations on parieto-temporal connections in the rhesus monkey. *Experimental Brain Research* 55, 301–312.

- Seltzer, B. and D. N. Pandya (1986). Posterior parietal projections to the intraparietal sulcus of the rhesus monkey. *Experimental Brain Research* 62, 459–469.
- Seltzer, B. and D. N. Pandya (1991). Post-rolandic cortical projections of the superior temporal sulcus in the rhesus monkey. *Journal of Comparative Neurology* 312, 625–640.
- Sereno, A. B. and J. H. R. Maunsell (1998). Shape selectivity in primate lateral intraparietal cortex. *Nature* 395, 500–503.
- Shadlen, M. N. and W. T. Newsome (1996). Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences USA* 93, 628–633.
- Shadlen, M. N. and W. T. Newsome (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience* 18(10), 3870–3896.
- Shelton, B. R. and C. L. Searle (1980). The influence of vision on the absolute identification of sound-source position. *Perception & Psychophysics* 28(9), 589–596.
- Shibutani, H., H. Sakata, and J. Hyvärinen (1984). Saccade and blinking evoked by microstimulation of the posterior parietal association cortex of the monkey. *Experimental Brain Research* 55, 1–8.
- Smith, D. R. and G. K. Smith (1965). A statistical analysis of the continual activity of single cortical neurones in the cat unanaesthetized isolated forebrain. *Biophysical Journal* 5, 47–74.
- Snowden, R. J., S. Treue, and R. A. Andersen (1992). The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Experimental Brain Research* 88, 389–400.
- Snyder, D. L. and M. I. Miller (1991). *Random point processes in time and space*. New York, NY: Springer-Verlag.
- Snyder, L. H., A. P. Batista, and R. A. Andersen (1997). Coding of intention in the posterior parietal cortex. *Nature* 386, 167–170.

- Snyder, L. H., A. P. Batista, and R. A. Andersen (1998). Change in motor plan, without a change in the spatial locus of attention, modulates activity in posterior parietal cortex. *Journal of Neurophysiology* 79(5), 2814–2819.
- Snyder, L. H., K. L. Grieve, P. Brotchie, and R. A. Andersen (1998). Separate body- and world-referenced representations of visual space in parietal cortex. *Nature* 394, 887–891.
- Softky, W. R. and C. Koch (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience* 13(1), 334–350.
- Sparks, D. L. and R. Hartwich-Young (1989). The deep layers of the superior colliculus. In R. H. Wurtz and M. E. Goldberg (Eds.), *The Neurobiology of Saccadic Eye Movements*, Chapter 5, pp. 213–255. Elsevier Science Publishers.
- Stanton, G. B., C. J. Bruce, and M. E. Goldberg (1995). Topography of projections to posterior cortical areas from the macaque frontal eye fields. *Journal of Comparative Neurology* 353, 291–305.
- Stein, J. F. (1989). Representation of egocentric space in the posterior parietal cortex. *Quarterly Journal of Experimental Physiology and Cognate Medical Sciences* 74, 583–606.
- Stricanne, B., R. A. Andersen, and P. Mazzoni (1996). Eye-centered, head-centered, and intermediate coding of remembered sound locations in area LIP. *Journal of Neurophysiology* 76(3), 2071–2076.
- Taira, M., S. Mine, A. P. Georgeopoulos, A. Murata, and H. Sakata (1990). Parietal cortex neurons of the monkey related to the visual guidance of hand movement. *Experimental Brain Research* 83, 29–36.
- Teich, M. C., R. G. Turcott, and R. M. Siegel (1996). Temporal correlations in cat striate-cortex neural spike trains. *IEEE Engineering in Medicine and Biology Magazine* 15(5), 79–87.
- Thier, P. and R. A. Andersen (1998). Electrical microstimulation distinguishes

- distinct saccade-related areas in the posterior parietal cortex. *Journal of Neurophysiology* 80, 1713–1735.
- Thurlow, W. R. and T. P. Kerr (1970). Effect of a moving visual environment on localization of sound. *American Journal of Psychology* 83, 112–118.
- Tolhurst, D. J., J. A. Movshon, and A. F. Dean (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research* 23(8), 775–785.
- Tolhurst, D. J., J. A. Movshon, and I. D. Thompson (1981). The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Experimental Brain Research* 41, 414–419.
- Tomko, G. J. and D. R. Crapper (1974). Neuronal variability: non-stationary responses to identical visual stimuli. *Brain Research* 79, 405–418.
- Tovée, M. J., E. T. Rolls, A. Treves, and R. P. Bellis (1993). Information encoding and the responses of single neurons in the primate temporal visual cortex. *Journal of Neurophysiology* 70(2), 640–654.
- Treves, A. and S. Panzeri (1995). The upward bias in measures of information derived from limited data samples. *Neural Computation* 7, 399–407.
- Ungerleider, L. G. and M. Mishkin (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior*, pp. 549–586. Cambridge, MA: MIT Press.
- Vaadia, E., D. A. Benson, R. D. Hienz, and M. H. Goldstein, Jr. (1986). Unit study of monkey frontal cortex: active localization of auditory and of visual stimuli. *Journal of Neurophysiology* 56(4), 934–952.
- Vallar, G., C. Guariglia, D. Nico, and E. Bisiach (1995). Spatial hemineglect in back space. *Brain* 118, 467–472.
- Victor, J. D. and K. P. Purpura (1996). Nature and precision of temporal coding in visual cortex: a metric-space analysis. *Journal of Neurophysiology* 76(2), 1310–1326.

- Victor, J. D. and K. P. Purpura (1997). Metric-space analysis of spike trains: theory, algorithms and application. *Network* 8, 127–164.
- Victor, J. D. and K. P. Purpura (1998). Spatial phase and the temporal structure of the response to gratings in V1. *Journal of Neurophysiology* 80, 554–571.
- Vogels, R., W. Spileers, and G. A. Orban (1989). The response variability of striate cortical neurons in the behaving monkey. *Experimental Brain Research* 77, 432–436.
- Vogt, C. and O. Vogt (1919). Allgemeine Ergebnisse unserer Hirnforschung. *Journal of Psychology and Neurology* 25, 279–462.
- Von Bonin, G. and P. Bailey (1947). *The Neocortex of Macaca mulatta*. London: University of Illinois Press.
- Wallace, M. T., L. K. Wilkinson, and B. E. Stein (1996). Representation and integration of multiple sensory inputs in primate superior colliculus. *Journal of Neurophysiology* 76(2), 1246–1266.
- Warren, D. H. (1970). Intermodality interactions in spatial localization. *Cognitive Psychology* 1, 114–133.
- Watanabe, M. (1992). Frontal units of the monkey coding the associative significance of visual and auditory stimuli. *Experimental Brain Research* 89, 233–247.
- Webster, D. B. (1992). An overview of mammalian auditory pathways with an emphasis on humans. In D. B. Webster, A. N. Popper, and R. R. Fay (Eds.), *The Mammalian Auditory Pathway: Neuroanatomy*, Chapter 1, pp. 1–22. New York, NY: Springer-Verlag.
- Webster, M. J., J. Backevalier, and L. G. Ungerleider (1994). Connections of inferior temporal areas TEO and TE with parietal and frontal cortex in macaque monkeys. *Cerebral Cortex* 5, 470–483.
- Welch, R. B. and D. H. Warren (1986). Intersensory interactions. In K. R. Boff, L. Kaufman, and L. P. Thomas (Eds.), *Handbook of Perception and Human Performance*, Chapter 25, pp. 1–36. New York, NY: Wiley Publishing Co.

- White, J. M., D. L. Sparks, and T. R. Stanford (1994). Saccades to remembered target locations: an analysis of systematic and variable errors. *Vision Research* 34(1), 79–92.
- Whittington, D. A., M.-C. Hepp-Raymond, and W. Flood (1981). Eye and head movements to auditory targets. *Experimental Brain Research* 41, 358–363.
- Xing, J., B. Stricanne, and R. A. Andersen (1994). A neural network model for sensorimotor transformations in macaque area LIP. *Society for Neuroscience Abstracts* 20, 143.
- Xu, L., S. Furukawa, and J. C. Middlebrooks (1998). Sensitivity to sound-source elevation in nontopographic auditory cortex. *Journal of Neurophysiology* 80, 882–894.
- Yao, L. and C. K. Peck (1997). Saccadic eye movements to visual and auditory targets. *Experimental Brain Research* 115, 25–34.
- Zambarbieri, D., G. Beltrami, and M. Versino (1995). Saccade latency toward auditory targets depends on the relative position of the sound source with respect to the eyes. *Vision Research* 23/24, 3305–3312.
- Zipser, D. and R. A. Andersen (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331, 679–684.