

UCSF

UC San Francisco Electronic Theses and Dissertations

Title

Explorations into the Dynamics of Oculomotor Tracking

Permalink

<https://escholarship.org/uc/item/190974k8>

Author

Schoppik, David

Publication Date

2007-06-27

Peer reviewed|Thesis/dissertation

Explorations into the Dynamics of Oculomotor Tracking
by
David Schoppik

DISSERTATION


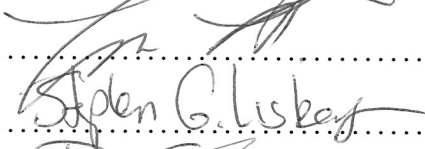



Submitted in partial satisfaction of the requirements for the degree of
DOCTOR OF PHILOSOPHY

in
NEUROSCIENCE

in the
GRADUATE DIVISION

of the
UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

Approved:

..... Michael S. Brainard, Ph.D.
..... Loren I. Frank, Ph.D.
..... Stephen G. Lisberger, Ph.D., Advisor
..... Philip N. Sabes, Ph.D., Chair
..... Krishna V. Shenoy, Ph.D., Outside Faculty
Committee in Charge

Deposited in the Library, University of California, San Francisco

.....
Date University Librarian

Degree Conferred:.....

UMI Number: 3274959

Copyright 2007 by
Schoppik, David

All rights reserved.

UMI[®]

UMI Microform 3274959

Copyright 2007 by ProQuest Information and Learning Company.
All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

Copyright 2007

David Schoppik

Acknowledgments

As I begin this work, I realize that I have no sufficient way to express my deep indebtedness to Katherine Nagel. She improves all that she touches: there is no aspect of this work, nor any part of my life that is not the better for her presence. I remain in awe, and thank her for her company, her thoughts, and her forbearance. Much of the intellectual foundation of this thesis comes from her, and I hope I've done those seminal ideas justice.

I am fortunate to have grown up in a family that valued learning, and I hope that this work does them justice. Without my grandmothers' courage, and my grandfathers' toil, I'd not be who I am today. Without my father's devotion, and my mother's nurturing independent thought from my earliest days, I'd surely not have made it this far. I've also got the best siblings anyone could ever hope for, who are living testimony of our parents' and grandparents' efforts. I can think of no better company that I've shared over the past few years than that of my friends, both within the neuroscience community and without.

Academically, the members of the Keck Center for Integrative Neuroscience at UCSF, particularly the members of my thesis committee have been more than generous with their time and patience. Without them and the atmosphere they create, this work could not have flourished as it did. What goes for the Keck Center goes triply so for the Lisberger lab, past and present, especially my office-mates in HSE-808. I've looked to you as a sounding board, and your echoes shaped this work. Finally, special mention is due Krishna Shenoy, who graciously agreed to participate in this thesis.

I'm at a loss for words as I search for those sufficient to describe Steve Lisberger's role

as my primary mentor at UCSF. Under his guidance, every aspect of my life as a scientist has been improved, and I will never be able to adequately express my gratitude for his time and energy over these past five years. Steve, may this thesis, and my career do you proud; so much of both will always be yours.

Finally, I dedicate this work to Juliet Olson. Her kindness and generosity have enabled me to lead an enviable life in the vibrant city that is San Francisco. Juliet's courage in the face of illness stands as inspiration and a constant reminder that neuroscience matters, not just to those of us seeking knowledge as an end in itself, but to all who strive for a better, healthier life. Juliet, while this work is too poor to provide any answers, please take comfort in knowing that it is the best I could do. Perhaps one day it may matter.

Abstract

All ethologically-relevant behavior is dynamic, and biological in origin. To understand the neural underpinnings of motor behavior is then to quantify the limitations placed on organisms as they move through and interact with their environment. This thesis stands as an attempt to understand the way the brain processes information to prepare for movement, and how a network of neurons might give rise to that movement. I chose to investigate oculomotor tracking in non-human primates as a model behavior that allowed excellent spatial and temporal control over the stimulus. First, by carefully controlling the motion of a single target, and measuring the eyes' response to particular perturbations, I mapped out effective constraints on the way the brain processes target motion to guide smooth eye movements. The next set of experiments was motivated by the observation that no two tracking movements are alike, even when faced with identical stimuli. I establish the dynamic relationship between fluctuations in single neurons recorded in the frontal eye fields and concomitant behavioral fluctuations. Utilizing measurements of the covariation between simultaneously recorded neurons, I describe the properties of the network of cells which give rise said fluctuations, which is likely dominated by fluctuations in a small subset of the active population at any moment in time. Taken together, the work I've done first establishes that proper oculomotor tracking likely takes the form of a re-allocation of cortical resources, but one that may be dominated by the activity of a small number of cortical neurons.

Contents

Acknowledgments	iii
Abstract	v
1 Spatial Processing of Motion	1
1.1 Abstract	1
1.2 Introduction	2
1.3 Materials and Methods	4
1.3.1 Subjects and equipment	4
1.3.2 Data collection and analysis	4
1.4 Results	7
1.4.1 A stimulus paradigm for controlling the distance between tracking target position and saccade endpoint	7
1.4.2 A spatial filter for processing motion direction around the endpoint of a saccade	11
1.4.3 Distracter targets do not affect the spatial filter	16
1.4.4 The spatial filter exists for saccades to stationary targets	19
1.4.5 Processing of target speed within the spatial filter	21
1.4.6 Frame of reference of the spatial filter for pursuit.	23
1.4.7 Dynamics of processing in the spatial filter for pursuit	26

1.4.8	Motion processing within the spatial filter reflects the weighted vector superposition of old and new target trajectories	28
1.5	Discussion	30
1.5.1	Old and new concepts of gain control for pursuit	30
1.5.2	Features of the neural underpinnings of a moveable spatial filter, and perceptual consequences	32
2	Sparse Encoding of Pursuit	35
2.1	Abstract	35
2.2	Introduction	36
2.3	Results	38
2.3.1	Neural and behavioral variation expressed as residuals	38
2.3.2	The structure of correlations between neural and behavioral variation	41
2.3.3	Filter shapes suggest an orthogonal basis set for movement commands	47
2.3.4	The timecourse of neuron-behavior correlations suggests sparse encoding, with low downstream noise	49
2.3.5	An example of the pairwise relationship between filtered spike trains	52
2.3.6	Quantification of the relationship between members of our population	55
2.3.7	Our model of the active population is compatible with a sparse set of neurons	58
2.3.8	Pairs of filters decorrelate spike trains	60
2.3.9	Control analyses	61
2.4	Discussion	62
2.4.1	Linear systems analysis of continuous oculomotor behavior, across the brain	62
2.4.2	A sparse network model in the face of co-variation	63
2.4.3	The role of spikes in the FEF_{SEM}	65

2.4.4	A sparse temporal code for movement variability in the cortex	66
2.4.5	The origin of behavioral fluctuations	67
2.4.6	What is a single spike worth?	68
2.5	Methods	68
2.5.1	Modeling the relationship between the residuals in spikes and eye movements	70
2.5.2	Data and analysis of paired recordings	73
3	Supplementary Material	76
3.1	Control Analyses	76
3.1.1	Abstract	76
3.1.2	Our findings obtain in the position domain, as well as for aligned data	78
3.1.3	Increasing bin size decreases our signal-to-noise ratio	80
3.1.4	A detailed evaluation of the predictions made by the example neuron	82
3.1.5	Generalization of neuron-behavior correlations to mean responses .	85
3.1.6	Generalization of neuron-behavior correlations across pursuit direction	88
3.1.7	Relationship between filter performance and neural response properties	89
3.2	Complementary Analyses	91
3.2.1	PCA decomposition of the filters	91
3.2.2	Evaluation of cross-validation fractions for filter generation and testing	92

List of Figures

1.1	Schematic diagram of the behavioral task and sample responses	8
1.2	Relationship between post-saccadic eye velocity and the distance from the endpoint of the saccade to the target position when the direction of new motion was orthogonal to the original motion	13
1.3	Relationship between post-saccadic eye velocity and the distance from the endpoint of the saccade to the target position in the presence of a distracter	18
1.4	Relationship between post-saccadic eye velocity and the distance from the endpoint of the saccade to the target position when the first saccade is aimed at a stationary target	20
1.5	Effect of the distance from the endpoint of the saccade to the target position on the relationship between post-saccadic eye velocity and target speed . . .	22
1.6	Analysis of the frame of reference of the relationship between post-saccadic eye velocity and the distance from the endpoint of the saccade to the target position	24
1.7	Temporal dynamics of the modulation of visual-motor transmission for pursuit	27
2.1	A typical example of neural and behavioral variability	40
2.2	The co-variation between spikes and eye movements for an example neuron	43
2.3	Properties of the filters that link neural activity to behavior	47

2.4	Temporal features of population activity, and a model of the consequences of downstream noise	50
2.5	The variation, co-variation, and neuron-behavior correlation for an example pair of simultaneously recorded neurons	53
2.6	A schematic network model, the relevant distributions from all pairs of cells, and a model implementation	56
2.7	A schematic network model, the relevant distributions from all pairs of cells, and a model implementation	58
2.8	The filters decorrelate paired spike trains	60
3.1	The relationship between spikes and behavior does not depend upon the representation of the eye movement, the trial alignment, or the trial length .	77
3.2	Smaller bin sizes lead to lower noise in the estimate of the relationship between spikes and eye movements	81
3.3	Predictions of the model relating spikes and behavior for the example neuron	83
3.4	Generalization from filters derived for pursuit in the preferred direction to the mean eye movement and the eye movement variation in different directions	87
3.5	The relationship between traditional neural response properties and the correlation between actual and predicted residual eye velocity	90
3.6	Principal component analysis decomposition of the filters	91
3.7	The effects on filter reliability of changing the ratio of trials used for filter generation to the fraction reserved for testing	93

Chapter 1

Saccades Exert Spatial Control of Motion Processing for Smooth Pursuit Eye Movements

1.1 Abstract

Saccades modulate the relationship between visual motion and smooth eye movement. Prior to a saccade, pursuit eye movements reflect a vector average of motion across the visual field. After a saccade, pursuit primarily reflects the motion of the target closest to the saccade's endpoint. We tested the hypothesis that the saccade produces a spatial weighting of motion around the endpoint of the saccade. Using a moving pursuit stimulus that stepped to a new spatial location just before a targeting saccade, we controlled the distance between the saccade's endpoint and the moving target's position. We demonstrate that the smooth eye velocity following the targeting saccade weights the pre-saccadic visual motion inputs by the distance from their location in space to the endpoint of the saccade, defining the extent of a spatiotemporal filter for driving the eyes. The center of the filter is located at the endpoint of the saccade in space, not at the position of the fovea. The filter

is stable in the face of a distracter target, is present for saccades to stationary and moving targets, and affects both the speed and direction of the post-saccadic eye movement. The spatial filter can explain the target-selecting gain change in post-saccadic pursuit, and has intriguing parallels to the process by which perceptual decisions about a restricted region of space are enhanced by attention. The effect of the spatial saccade plan on the pursuit response to a given retinal motion describes the dynamics of a coordinate transformation.

1.2 Introduction

Retinal inputs provide sensory drive for visually guided behaviors. Because the eyes move continuously, the retinal locus of a stimulus cannot provide unambiguous guidance about the correct location of objects in space. Yet, spatial information drives many behaviors. For example, saccadic eye movements are based on the location of targets in space, not on the retinal locus activated by a momentary flash [Mays and Sparks, 1980]. Transformation of inputs that enter the brain in retinal coordinates into spatial behavior is a fundamental sensorimotor process. Saccadic selection of moving targets for smooth pursuit eye movements provides a specific example of this transformation that lends itself particularly well to study.

Recent reports have established that visual motion processing for smooth pursuit eye movements can be altered dynamically by saccades. An example occurs when a single target moves across a locus eccentric from the fovea: the initial pursuit response is weak, but the response after the saccade is strong, as if the visual motion drive for pursuit has been enhanced. Similarly, if two moving targets are present, the saccade to one is an oculomotor choice: post-saccadic smooth eye velocity is dominated by visual inputs from the moving target at the endpoint of the saccade, even if the initial pursuit was driven by the other target [Gardner and Lisberger, 2001, 2002, Liston and Krauzlis, 2005]. These results illustrate the ability of the spatially specific saccade plan to enhance motion inputs that

arise in a retinal coordinate frame. The preferential treatment of one target's motion at the expense of the others resembles cued spatial attention.

Three features of perceptual attention parallel the effects of saccades on subsequent pursuit. First, both specify a region of space for enhanced processing: saccades determine the location of target motion to be preferentially processed by pursuit while spatial attention enhances perception at a region of space. Second, both can be triggered by the saccade plan. For pursuit, even sub-threshold stimulation in the superior colliculus (SC) biases the subsequent pursuit choice towards targets in the movement field of the excited neurons [Carello and Krauzlis, 2004]. For perception, stimulation of the frontal eye fields (FEF) with current too weak to evoke a saccade is sufficient to cause the responses of extrastriate neurons in area V4 to mimic the effects of cued attention [Moore and Armstrong, 2003]. Similarly, microstimulation in the SC enhances attentional processing in a spatially selective manner [Muller et al., 2005]. Third, both depend on modulation of response gain. Target choice for pursuit reflects in an increase in the response to specific visual motion inputs. Likewise, attention manifests itself as increases in the gain of neural responses for stimuli falling within a restricted spatial locus. If the behavioral effects of target choice for pursuit and perceptual attention are so similar, then perhaps they use similar neural mechanisms.

We tested the hypothesis that saccades enhance the processing of visual motion within a specific spatial location, just as does cued attention. We show that the gain of the pursuit response to a given target motion is weighted spatially, by the distance from the endpoint of the upcoming saccade to the spatial location of the target. We propose that a spatial filter anticipating the movement of the eyes controls the gain of visual-motor transformation for pursuit.

1.3 Materials and Methods

1.3.1 Subjects and equipment

Four male monkeys (*Macaca mulatta*) were used in experiments approved by the Institutional Animal Care and Use Committee of the University of California, San Francisco. All experimental procedures were in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals. Briefly, eye position was monitored using the scleral search coil technique while the head was held stationary using custom hardware that had been implanted on the monkey's skull during sterile surgery under Isoflurane anesthesia. Surgical procedures have been reported before [Ramachandran and Lisberger, 2005]. All monkeys were experienced in smoothly tracking visual targets for a fluid reward, and needed no additional behavioral training for this study.

In a dimly lit room, visual stimuli were projected onto the back of a tangent screen placed 114 cm in front of the monkey. Projecting a stationary red light-emitting diode onto the center of the screen generated a fixation spot. The position of the tracking target, which was a small white 0.5-deg spot, was controlled by reflecting a beam of light off a pair of two orthogonal mirror galvanometers. Many of the experiments described here implicitly rely on displacing the target instantaneously; the galvanometers could complete our largest step (9 deg) within 4 ms. To verify that our data were not affected by the visual streak generated during the displacement, we also presented targets using two separate pairs of galvanometers and shutters to illuminate different targets before and after displacement, eliminating the visual streak. We did not find any differences in the results for the two methods of target presentation, and data from both are included.

1.3.2 Data collection and analysis

Stimuli were presented as a series of trials, where each trial represented a different set of target trajectories, and the order of the trials was randomized. Each trial began when the

fixation target appeared for a random time (500-1100 ms), followed by the appearance of the tracking target. The target usually appeared slightly eccentric (0-1.5 deg) to the position of fixation and immediately moved toward the position of fixation. On half of the trials, the target moved for a short interval (100-200 ms) and then underwent a step change in position, sometimes changing its direction of motion or speed as well. Whether or not the target underwent this step, it continued moving for a random duration (600-800 ms) during which the monkey was required to keep his gaze within 2 deg of the tracking target. Whenever the target was displaced, the monkey was allowed a 300-ms grace period when fixation requirements were suspended so that he could catch the target and complete the trial. In each experiment, all variables such as the initial and post-displacement direction of motion were balanced to discourage any sort of learning, as well as efforts to anticipate the target motion. Monkeys were rewarded with droplets of fluid for the approximately 95% of all attempted trials they completed. They typically worked until satiated, for approximately 1500-2500 trials.

Horizontal and vertical eye position signals were sampled at 1 kHz on each channel and differentiated digitally using a balanced difference algorithm over a 2 ms time window. The differentiated traces then were filtered digitally using a 2-pole Butterworth filter with a cutoff at 25 Hz. For a discussion of the effects of filtering on the measurement of post-saccadic eye velocity, see [Lisberger, 1998]. An automated threshold-crossing algorithm marked the beginning and end of each saccade with a threshold at 50 deg/s; the ends of the saccades were provisionally estimated to occur 15 ms after eye velocity crossed back below threshold. For target motion at 10 deg/s, we lowered the threshold to 40 deg/s to catch smaller saccades. After provisional marking of saccade start and end by the automated algorithm, all trials were inspected and, if necessary, the saccade marks were adjusted by hand using custom software. The measurements reported in this paper rest on a precise identification of saccade initiation; to confirm that our algorithm gave excellent results, we re-marked the saccades by hand for selected experiments on the basis of

the initial deflection on eye velocity traces that had been generated on-line using an analog circuit. The circuit differentiated signals up to a cutoff at 25 Hz and filtered signals of higher frequencies. Because the analog filter was causal, the latter assessment of the time of the start of the saccade is uncorrupted by the smoothing that arises with zero-phase digital filtering. Results did not depend on how the saccade start and end were determined.

Our analysis used standard statistical methodology, including t-tests and linear regression techniques. Data were fit to simple linear models in a least-squared sense, and the p-values for slopes and y-intercepts were estimated by generating a t-statistic for each. We fitted plots of post-saccadic eye velocity as a function of the duration of the motion that drove it with either a sigmoid (1) or an exponential (2):

$$\dot{E}(t) = \frac{1}{1 + e^{-m(t-\beta)}} \quad (1.1)$$

$$\dot{E}(t) = e^{\beta + (\tau * t)} \quad (1.2)$$

For these equations, m represents the slope of the function, τ the time constant, and β the point at which it reaches half its maximum value. When fitting the sigmoid, we used only eye velocity data that had been normalized to the mean eye velocity on control trials; this removed our subjects' spatial bias from the analysis. Curve-fitting was done in a least-squares sense, using Levenberg-Marquardt minimization; 95% confidence intervals were estimated from the residuals. To find the time at which the sigmoid was at 10% and 90% of its maximal value, we solved (1.1) for t . The original data used in this paper, analysis code, and figure-generation code are available on the internet (<http://keck.ucsf.edu/~schoppik/>).

1.4 Results

1.4.1 A stimulus paradigm for controlling the distance between tracking target position and saccade endpoint

Our experiments tested the hypothesis that saccades enhance the processing of visual motion inputs for pursuit in proportion to the proximity of the target to the endpoint of the saccade. We began by designing a behavioral task that would allow us to control the location of a moving target relative to the endpoint of a saccade, by disrupting the usual accuracy of saccades to moving targets. Specifically, we contrived for monkeys to make a saccade to a target that was no longer where they expected it to be, but instead had been displaced some distance at carefully controlled moments. Our approach bears some resemblance to the well-explored double-step paradigm [Becker and Jurgens, 1979], and to remapping experiments in the parietal cortex [Duhamel et al., 1992].

Our task, which can be described in four intervals, is diagrammed in two ways in Figure 1.1. First, the set of four Cartesian plots in 1.1A.1 - 1.1A.4 provide snapshots of both target and eye motion for each interval. Second, for an example trial, Figure 1.1B plots the position of the eyes and the target as a function of time, and Figure 1.1C similarly plots velocity. The black bars with numbers over them delineate the duration of each of the four intervals specified in Figures 1.1A.1 - 1.1A.4, allowing the reader to cross-reference between the spatial and temporal presentations.

Each trial begins with a brief period of fixation, after which the target moved rightward (Figure 1.1A.1). The initial motion of the target was designed both to necessitate a rightward saccade and to drive the initiation of pursuit. The short-latency pursuit response is shown as a small black arrow pointing rightward in Figure 1.1A.2 (second interval). For the duration of the second interval, the initial saccade plan is locked in; the animal is now committed to use the target position at the end of interval 1 to drive a rightward saccade. During the second interval, though, the target's motion changes: it is displaced vertically

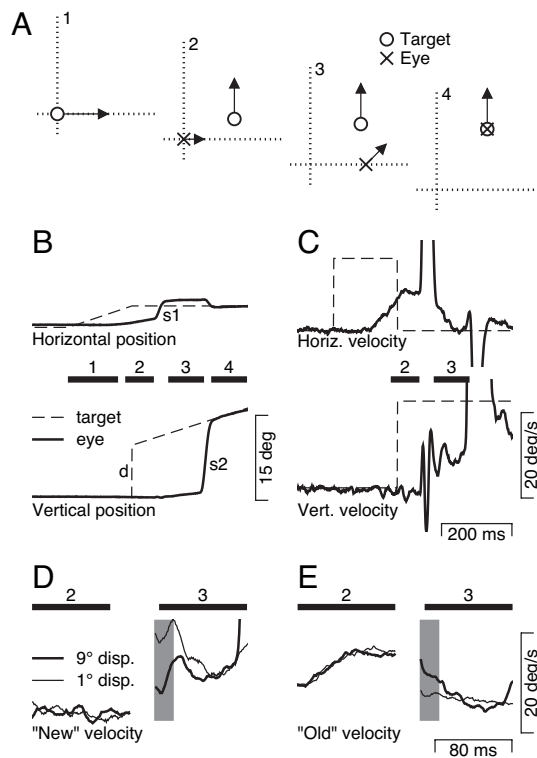


Figure 1.1: Schematic diagram of the behavioral task and sample responses. A: The task is shown in spatial coordinates as a sequence of 4 snapshots. In each panel, the open circle and the cross indicate the position of the target and eye. A1. Starting just off fixation point, the target moves to the right at 20 deg/s. A2. The eye moves rightward in response to the initial target motion, while the target stops and begins moving upwards at 20 deg/s (the "new" motion). A3. The eye makes a rightward saccade driven by the pre-displacement location of the target, while the target continues to move up. A4. A second saccade takes the eye directly to the target, which is tracked it appropriately. B: The horizontal and vertical components of eye (solid line) and target (dashed line) positions are plotted as a function of time. Black bars with numbers over them mark the periods corresponding to the schematic in (A). The labels "d", "s1", and "s2" mark the target displacement to the new position and motion, and the first and second saccades. C: The horizontal and vertical components of eye velocity (solid line) and target velocity (dashed line) as a function of time. The bold, horizontal bars correspond to intervals 2 and 3. D, E: Same data as in C, but with the time base expanded so that data are shown only for intervals 2 and 3. In D-E, the thick and thin lines show responses when the size of the target displacement was 9 deg and 1 deg, respectively, while the shaded vertical rectangles indicate the interval used for analysis.

0, 1, 5, or 9 deg and its direction of motion changes from horizontal to vertical. In the third interval (Figure 1.1A.3), we see the consequences of the saccade planned to the target position in interval 1. The rightward saccade takes the eyes to the "wrong" location; however, the post-saccadic pursuit, which must be driven by pre-saccadic target motion, has already incorporated both the "old" target motion of interval 1 and the "new" target motion of interval 2 (diagrammed by an arrow moving both up and to the right). By analyzing the component of eye motion in the direction of the "new" target motion, as a function of the jump size, we test the hypothesis that the target motion affects pursuit eye velocity with a gain related to the location of the new target motion in space. Finally, in interval 4 (Figure 1.1A.4), the monkey makes a corrective saccade to the new location of the target, tracking accurately for a brief period of time to receive his reward.

Figure 1.1B diagrams the eye and target positions as a function of time, while 1C similarly diagrams the corresponding velocities. The initial horizontal target motion marks the beginning of interval 1, and the eye's horizontal smooth response can be seen slightly before interval 2. The vertical target displacement and subsequent motion (marked 'd') marks the beginning of interval 2. The initial saccade (marked 's1'), driven by the horizontal target trajectory, defines the end of interval 2. Before the beginning of interval 3, the change in eye position is entirely horizontal, bringing the eye ahead of the target. At the end of saccade s1, interval 3 begins; the eye velocity traces reveal that despite the switch from horizontal to vertical target velocity, both horizontal and vertical eye velocity is present up until the point of the second, corrective saccade (marked 's2'), which marks the start of interval 4.

The success of our experiment depends on correctly timing the length of interval #1. The endpoint of saccade s1 must be determined by the original position of the moving target, but the saccade latency must be long enough for pursuit to have seen the new target motion. Because pursuit has a shorter latency to respond to the new target motion, we can hope that interval 2 will interpose up to 100 ms of new target motion between displacement

d and saccade s_1 . As a consequence, post-saccadic eye velocity is driven by the new target motion and is our dependent variable. Furthermore, because we displaced the target at the start of interval 2, the new motion is some distance above the endpoint of the first saccade and so the distance from the endpoint of the saccade to the target becomes our independent variable. To ensure that the monkeys did not simply withhold their saccades until the target had changed direction, 50% of trials were controls in which the target did not change its trajectory, and the monkey received his reward by tracking accurately to the end of the trial.

On 45% of the trials, the monkey made his first saccade before the target had changed direction (interval #1 was too long) or made a saccade directly to the post-displacement location of the target (interval #1 was too short); both outcomes are to be expected when the target trajectory changes in this fashion [Lisberger et al., 1975]. Such trials were excluded from further analysis, as the pre-saccadic motion took place too close to the spatial endpoint of the saccade. After excluding control trials (50% of all trials) and those with unsuitable saccade timing (45-48% of all trials), we were left with usable trials that represented 2-5% of the all trials the monkeys performed. For the analyses in Figures 2-5, we further restricted the trials we analyzed to those in which the duration of the new motion was greater than 50 ms. As a result, many days of experiments were needed to obtain sufficient data.

Finally, although it may appear as if the second saccade has a shorter latency than the first saccade, when measured relative to the onset of the driving target motion, they are quite similar. The average delay from the time the target changed position and direction to this second saccade ($d \rightarrow s_2$) was 141 ms (± 28 ms SD). The average delay from the onset of target motion to the first saccade was similar at 146 ms (± 32 ms). The time between the onsets of the first and second saccades ($S_1 \rightarrow S_2$) averaged 82 ms (± 36 ms). As would be expected in a double-step-like task, the endpoints of the saccade were not influenced by the ongoing pursuit, and had a bimodal distribution, landing either directly on the target,

or at the "incorrect" location.

1.4.2 A spatial filter for processing motion direction around the endpoint of a saccade

The expanded eye velocity traces in Figures 1.1D and E show examples of how the distance between the moving target and the endpoint of the saccade itself affects post-saccadic eye velocity. Each panel shows the eye velocity along one component, for two trials: one with a 1 deg step (thin traces) and one with a 9 deg step (thick traces). Figure 1.1D illustrates the eye velocity responses along the axis of new motion, here vertical, and Figure 1.1E the eye velocity responses along the axes of old motion, here horizontal. Compared to when the step was 1 deg (thin traces), the vertical, or new, component of post-saccadic eye velocity (Figure 1D) at the start of interval #3 was smaller and the horizontal, or old, component (Figure 1.1E) was larger when the target had been displaced 9 deg (bold traces) from the endpoint of the saccade. Consequently, pursuit was more nearly vertical when the endpoint of the saccade was close to the vertically moving target. Alternately, when the endpoint of the saccade was far from the vertical motion, pursuit was more nearly horizontal.

The records in Figure 1.1D and E present an example of our basic finding, and also underscore the challenges of performing and analyzing these experiments. First, both the horizontal and vertical eye velocity traces in the time between interval #2 and interval #3 are contaminated by rapid deflections associated with the saccade to the prior location of the original rightward-moving target. We have previously provided a detailed evaluation of the challenges of measuring immediate post-saccadic eye velocity and we are confident that the differences demonstrated in Figures 1D and E are well beyond the time when eye velocity is related directly to the saccades themselves [Lisberger, 1998]. Second, because the feedback loop closes quickly after the end of the saccade, by design, our analysis is based on a brief 20 ms interval (shaded rectangle in Figures 1.1D and E) when eye velocity

depends on the distance from the endpoint of the saccade to the moving target. As the pursuit system naturally corrects the difference between target and eye velocity, the eye velocity traces for target displacements of 1 and 9 deg converge quickly, even though they are quite different during the period under study.

Figure 1.2B uses data from monkey Cb to illustrate how we quantified the results of our experiments. Each of the 1672 points represents the post-saccadic pursuit velocity in a single trial plotted as a function of the radial distance in space between the target and the eye at the endpoint of the saccade. Variation in the latency and accuracy of the monkey's first saccade cause scatter along the x-axis for each cluster of points and variation in the post-saccadic eye velocity cause scatter along the y-axis. The experiment summarized in Figure 2B contained targets that began to move along the horizontal meridian. In 50% of the trials, targets then stepped 1, 5, or 9 deg vertically and moved either upward or downward, shown as points plotted at negative or positive values along the y-axis. Both the clouds of points from the individual trials and the averages for monkey Cb (open diamonds) show that the post-saccadic vertical eye velocity depended strongly on the distance from the endpoint of the saccade to the spatial location of the new motion. Trials were included in this analysis only if the first saccade landed within 1 deg of the horizontal meridian, and if the interval of new motion was at least 50 ms long. Requiring a longer period of target motion before the saccade would have revealed larger effects, but at the cost of including fewer trials, a feature of the data that we will document later.

To summarize the results, we inverted the eye velocity traces from the trials that ended with downward motion, so that we could combine the data from both directions of target motion as if the new motion always had been upward. We then averaged the data for each displacement of target position (0/1, 5, or 9 deg) and plotted the mean and standard deviation of post-saccadic eye velocity as a function of the mean distance from the target position to the eye position at the endpoint of the saccade (Figure 1.2C). For all 4 monkeys, the post-saccadic eye velocity in the direction of the new target motion was largest when

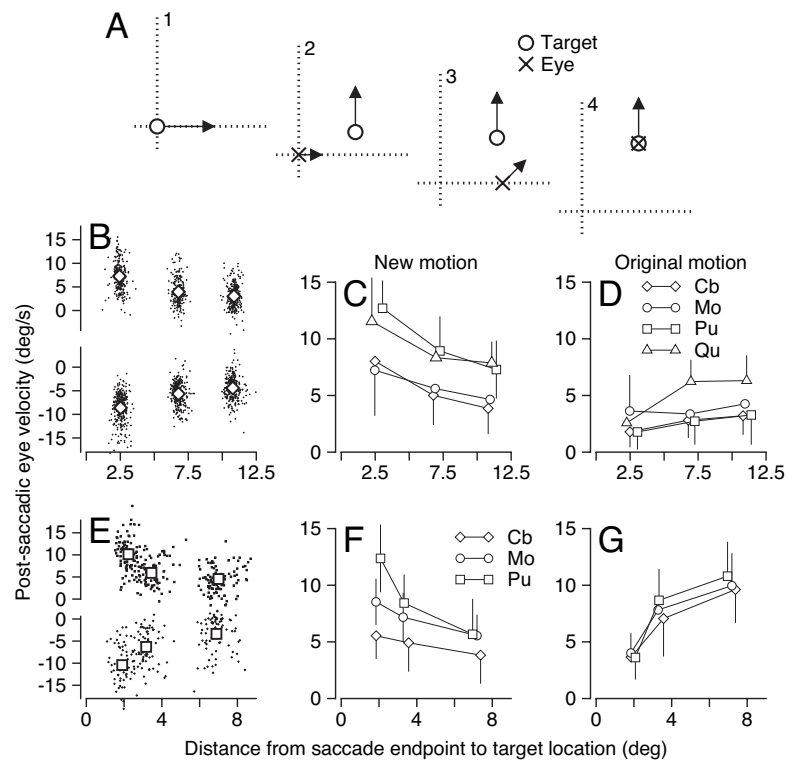


Figure 1.2: Relationship between post-saccadic eye velocity and the distance from the endpoint of the saccade to the target position when the direction of new motion was orthogonal to the original motion. A: Same schematic as Figure 1.1. B-D: Analysis of data when target displacement was orthogonal to the original motion, as shown schematically in A. E-G: Analysis of data when target displacement was parallel to the original motion. B, E: Plots of post-saccadic eye velocity in the direction of new motion as a function the distance between target and eye at the end of the saccade for individual monkeys. Dots show data for individual trials and larger symbols (diamonds in B for monkey Cb and squares in E for monkey Pu) show means for the 3 different amplitudes of displacement of target position. Positive and negative values of post-saccadic eye velocity were taken from trials in which the new target motion was upward or downward. C, F: Means and standard deviations of post-saccadic eye velocity in the direction of new motion for all monkeys used in each experiment. Different symbols correspond to different monkeys. E, G: The mean and standard deviation of post-saccadic eye velocity in the direction of the original motion. The graphs in C and D summarize data for experiments in which the original direction of motion was horizontal and new motion vertical, while in F and G, original motion was vertical and new motion horizontal. Before averaging, we took the absolute velocity of all eye velocities.

the distance between the target and eye position at the endpoint of the saccade was smallest. Post-saccadic eye velocity declined as the distance increased. In the same trials, the eye velocity response to the original, horizontal target motion showed the opposite trend (Figure 1.2D); the greater the distance from target to saccadic endpoint, the faster the eyes moved along the original component of motion.

Statistical evaluation supported the qualitative effects seen in Figures 1.2B-D. Regression lines were fit to the collection of single-trial measurements from each monkey individually; complete statistical information can be found in Table 1.4.2. The mean slope across monkeys was 0.45 ± 0.14 (SD), and the mean y-intercept was 10.72 ± 2.88 (SD). All slopes and y-intercepts were significantly different from zero ($p < .001$). In Figure 1.2D, the mean slope was 0.35 ± 0.20 (SD), but was significantly different from zero for only two monkeys. The absence of a strong effect in Figure 1.2D is due to relationships we will document later between post-saccadic eye velocity in the original and new directions of motion and the duration of the interval of new motion.

We repeated the experiment in Figure 1.2A-D, with the motion rotated so that the original direction was vertical, and the new direction horizontal (data not shown, details for each monkey in Table 1.4.2). Statistical evaluation demonstrated the same decrease in post-saccadic eye velocity with increasing distance between the target and the endpoint of the saccade. For four monkeys, neither the mean slope in the new direction, -0.14 ± 0.28 (SD), nor the mean y-intercept, 8.44 ± 2.97 (SD) were different from the data shown in Figure 2C (paired t-test, $p > 0.06$, $p > 0.08$). Similarly, neither the mean slope in the old direction, 0.20 ± 0.14 (SD), nor the mean y-intercept 2.83 ± 0.60 (SD), were different from the data shown in Figure 2D (paired t-test, $p > 0.36$, $p > 0.42$).

To control for the effect of step direction, we also conducted a variation of the experiment in Figure 1.2A-D where the step of target motion at the start of interval #2 was parallel to the initial motion of the target, rather than orthogonal. Otherwise, the experimental design was the same: at the time of the target step, the direction of smooth target motion

changed so that it was orthogonal to the original motion, and the amplitudes of the target stepdf again were 1, 5, or 9 deg. The step always moved the target away from the position of fixation, and the experiment was run only with target motion that was initially vertical and then changed to new motion that was horizontal.

For an example monkey (Pu), the cloud of points of post-saccadic eye velocity from individual trials ($n=285$) and the averages (Figure 1.2E, squares) indicate that the speed of eye motion in the direction of the new target motion was related inversely to the distance between target and eye position at the endpoint of the saccade. For all 3 monkeys tested on this experiment (Figure 1.2F), the same trend is visible in the plots of average post-saccadic eye velocity as a function of the distance from the endpoint of the saccade to the target. Further, for all 3 monkeys, the post-saccadic eye velocity in the direction of the original target motion scaled proportionately with the distance to the endpoint of the saccade (Figure 1.2G). The fact that the new direction of motion was horizontal, the monkeys' best axis of pursuit, led to a better illustration of this feature of our data in Figure 1.2G than in Figure 1.2D. Statistical analysis supported the trends seen in the graphs. For eye velocity in the new direction of target motion, the mean slope of the three regression lines fitted to all trials from each monkey independently was -0.66 ± 0.41 (SD), and the mean y-intercept was 9.64 ± 3.67 (SD). Neither the slopes nor the y-intercepts for Figure 1.2F were significantly different from those for Figure 1.2C (two-tailed paired t-test, $p > 0.4$). For eye velocity in the original direction of target motion (Figure 1.2G), the mean slope of the relationship between post-saccadic eye speed and distance from saccade endpoint to target was 1.06 ± 0.06 , and the mean y-intercept 2.84 ± 0.33 . Here, unlike Figure 1.2D, the slopes and y-intercepts were significant for all three monkeys tested ($p < 0.001$).

The data in Figure 1.2 show that the pursuit response to a given target motion is modulated by the distance from the endpoint of the saccade to the target. The modulation is consistent with the conclusion that a spatial filter surrounds the endpoint of a saccade. If the target is near the center of the spatial filter, its motion has priority access to the pursuit

system and dominates the smooth eye velocity response.

1.4.3 Distracter targets do not affect the spatial filter

In experiments like that diagrammed in Figure 1.3A, interval #1 provided two targets moving in opposite directions along the vertical axis. The monkey was free to saccade to either target, and each had an equal likelihood of undergoing a displacement and a change to a new, horizontal direction of motion in interval #2, just as in the experiments of Figure 2. Software detected saccades, determined which target the monkey had chosen, and turned off the remaining target, allowing the monkey to track only one target and receive a reward. Trials were not included in the analysis if the monkey chose the target that did not displace (25% of all trials), or if neither target was displaced (50% of all trials). When the monkey chose the target that did displace, his behavior was largely the same as during the single target case, in that his saccade was too early or late most of the time (23% of all trials), leaving us with only 2% of all trials for analysis.

Figure 1.3B shows the results of the experiment for one monkey (Qu). Each of the 181 points shows the pursuit response on a single trial. The y-axis plots the horizontal component of post-saccadic eye velocity when the step of target position in interval #2 was to the right or left (points plotting at positive or negative values on the y-axis); the x-axis shows the radial distance between target and eye at the end of the saccade. Both the clouds of points for the individual trials and their means (open triangles) indicate that the post-saccadic horizontal eye speed tends to be slower when the distance from the endpoint of the saccade to the horizontally-moving target is larger. The same relationship appears in the means of post-saccadic horizontal eye velocity of all four monkeys (Figure 1.3C). Regression lines were fit to all trials from each monkey separately; detailed statistics are in Table 1.4.2. The regression slope of eye velocity versus distance for four monkeys averaged -0.55 ± 0.24 , the y-intercept averaged 10.42 ± 4.04 , and neither the slope nor the y-intercept was statistically different from those for the two experiments summarized in Figure 1.2

Table 1.1: Parameters of linear regression for text and Figures 1.2-1.4. In the left column, the first entry indicates the monkey. When the data appeared in a figure, the second entry indicates the figure panel. When the data was mentioned only in the text, the second entry indicates whether the regression is for the old or new direction of motion and the third entry is "text."

Monkey, Figure	# Trials	Slope (95% CI)	y-intercept (95% CI)	R^2	F	p
CB, 2C	1672	-0.50 (0.04)	8.96 (0.29)	0.28	664.81	0
MO, 2C	459	-0.28 (0.08)	7.74 (0.69)	0.09	44.85	0
PU, 2C	42	-0.61 (0.26)	14.02 (2.13)	0.36	22.84	0
QU, 2C	29	-0.41 (0.30)	12.15 (2.29)	0.22	7.71	0
CB, 2D	1672	0.17 (0.02)	1.48 (0.16)	0.12	230.76	0
MO, 2D	459	0.06 (0.09)	3.41 (0.71)	0	1.63	0.20
PU, 2D	42	0.15 (0.20)	1.55 (1.66)	0.05	2.21	0.14
QU, 2D	29	0.40 (0.20)	2.10 (1.57)	0.38	16.40	0
CB new, tex	1295	-0.18 (0.03)	5.48 (0.25)	0.08	111.31	0
MO new, text	621	-0.10 (0.05)	7.59 (0.40)	0.03	17.76	0
PU new, text	117	-0.47 (0.18)	12.55 (1.43)	0.19	26.94	0
QU new, text	295	0.21 (0.09)	8.14 (0.64)	0.07	21.72	0
CB old, text	1295	0.04 (0.40)	3.01 (0.30)	0.27	378.05	0
MO old, text	621	0.46 (0.56)	1.95 (0.48)	0.27	230.14	0
PU old, text	117	0.47 (0.13)	3.30 (1.06)	0.30	49.91	0
QU old, text	295	0.37 (0.08)	3.06 (0.59)	0.21	81.87	0
CB, 2F	188	-0.34 (0.13)	6.21 (0.66)	0.12	25.65	0
MO, 2F	138	-0.52 (0.14)	9.20 (0.64)	0.28	52.32	0
PU, 2F	285	-1.13 (0.81)	13.51 (0.81)	0.39	184.40	0
CB, 2G	188	1.02 (0.15)	2.46 (0.76)	0.48	173.06	0
MO, 2G	138	1.04 (0.16)	3.04 (0.73)	0.54	158.95	0
PU, 2G	285	1.12 (0.16)	3.03 (0.81)	0.39	182.00	0
CB, 3C	278	-0.22 (0.07)	5.39 (0.52)	0.12	37.69	0
MO, 3C	141	-0.52 (0.13)	9.00 (0.95)	0.31	61.20	0
PU, 3C	286	-0.66 (0.10)	13.00 (0.80)	0.35	154.62	0
QU, 3C	181	-0.78 (0.13)	14.30 (1.05)	0.45	149	0
CB, 3D	278	-0.15 (0.07)	3.88 (0.47)	0.07	21.56	0
MO, 3D	141	0.24 (0.10)	1.92 (0.74)	0.13	21.61	0
PU, 3D	286	0.01 (0.06)	2.48 (0.49)	0	0.21	0.64
QU, 3D	181	0.03 (0.09)	2.46 (0.73)	0	0.44	0.51
CB, 4C	248	0 (0.11)	4.45 (0.91)	0	0	0.96
MO, 4C	780	-0.35 (0.06)	7.60 (0.50)	0.15	133.23	0
PU, 4C	211	-0.61 (0.12)	10.41 (1.02)	0.31	93.92	0
QU, 4C	615	-0.41 (0.07)	9.34 (0.56)	0.17	123.38	0

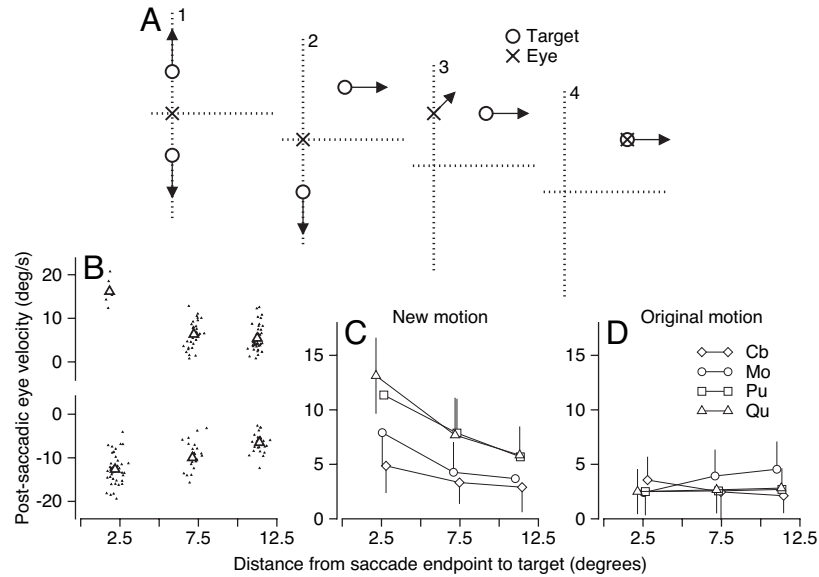


Figure 1.3: Relationship between post-saccadic eye velocity and the distance from the endpoint of the saccade to the target position in the presence of a distracter. A: Schematic diagram showing the task as a sequence of 4 snapshots in time. A1. Two targets (open circles) began moving in opposite directions away from the fixation point and the eye (x). A2. One of the targets is displaced and begins to move orthogonally, while the other continues along the original trajectory. A3. A saccade is made to the old location of one of the targets and the eyes begin to move smoothly; the second target disappears. A4. A saccade was made to the new location of the target, followed by proper tracking. B: Plot of post-saccadic eye velocity in the direction of new motion as a function the distance between target and eye at the end of the saccade for one monkey. Dots show data for individual trials and larger symbols show means for the 3 different amplitudes of displacement of target position. Positive and negative values of post-saccadic eye velocity were taken from trials in which the new target motion was rightward or leftward. C: Means and standard deviations of the absolute value of post-saccadic eye velocity in the direction of new motion for all monkeys used in each experiment. Different symbols correspond to different monkeys. D: The mean and standard deviation of the absolute value of post-saccadic eye velocity in the direction of the original motion.

(two-tailed paired t-test $p > 0.3$). Regression analysis was performed on all trials that went into Figure 1.3D, and detailed statistics can be found in Table 1. As in Figure 1.2D, the relationship between vertical post-saccadic eye velocity and the distance from the eye to the target at the endpoint of the saccade is unimpressive (Figure 1.3D), though the slopes for two of the monkeys are statistically different from zero.

1.4.4 The spatial filter exists for saccades to stationary targets

The experiment diagrammed in Figure 1.4A follows same structure used before, except that in interval 1, the target does not move smoothly. Each trial began with the monkey fixating straight ahead. At the start of interval #1, the fixation target disappeared and another target appeared 3 degrees eccentrically along the horizontal meridian, to provide the stimulus for a saccade. In interval 2, the target stopped vertically and begins to move horizontally. The target remained stationary in 50% of the trials to encourage the animal to make a timely and accurate saccade; these control trials were removed from further analysis.

As before, a falloff in post-saccadic eye velocity as a consequence of the distance between target and saccadic endpoint appears in the clouds of points for individual trials for one monkey (Figure 1.4B), the means for those trials (open squares in Figure 1.4B), and the means for all 4 monkeys used in this experiment (Figure 1.4C). Regression lines were fit to data from all trials for each monkey separately, and detailed statistics can be found in Table 1.4.2. Coefficients for one monkey (Cb) were not significantly different from zero, indicating that a regression model was not appropriate. As such, his data were excluded from the following summary statistics. The average regression slope for the remaining three monkeys in Figure 4C was -0.46 ± 0.14 , and the y-intercept was 9.12 ± 1.42 . Neither the slope nor the y-intercepts were significantly different from those for the experiments summarized in Figures 1.2 and 1.4 (two-tailed paired t-test, $p > 0.15$). For this experiment it did not make sense to show plots of the orthogonal, vertical eye velocity as there was

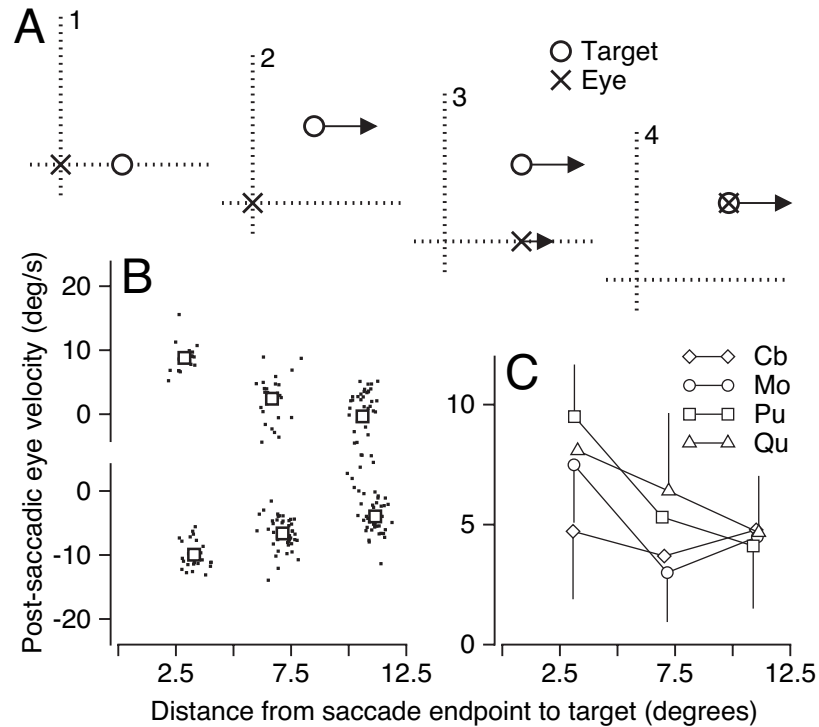


Figure 1.4: Relationship between post-saccadic eye velocity and the distance from the endpoint of the saccade to the target position when the first saccade is aimed at a stationary target. A: Schematic diagram showing the task as a sequence of 4 snapshots in time. A1. A target (open circle) appears stationary but offset from the current position of the eye (x). A2. The target suddenly moves to the right at 20 deg/s. A3. A saccade is made to the old location of the target and the eyes begin to move smoothly along the direction of target motion. A4. The eye saccades to the new location of the target and tracks it appropriately. B: Plot of post-saccadic eye velocity in the direction of new motion as a function of the distance between target and eye at the end of the saccade for one monkey. Dots show data for individual trials and larger symbols show means for the 3 different amplitudes of displacement of target position. Positive and negative values of post-saccadic eye velocity were taken from trials in which the new target motion was rightward or leftward. C: Means and standard deviations of the absolute value of post-saccadic eye velocity in the direction of new motion for all monkeys used in each experiment. Different symbols correspond to different monkeys.

never any vertical target motion.

1.4.5 Processing of target speed within the spatial filter

The experiment diagrammed in Figure 1.5A has the same basic structure as the previous experiments. In interval #1, the target started from straight ahead gaze and moved to the right. At the start of interval #2, after the initiation of pursuit, but before the monkey executed a catch-up saccade for the initial target motion, the target stepped to a location that was 1, 5, or 9 deg off the horizontal meridian and continued to move in the same direction. At the time of the target step, the target either continued at its original speed of 20 deg/s (Figure 1.4A, white circle), or changed to 10 or 30 deg/s (black and grey circles, respectively). As a result, the new motion comprised one of the three target speeds that was present in interval #2, before the catch-up saccade; the speed of new motion should affect post-saccadic eye velocity more strongly when the distance from the target to the eye position at the endpoint of the saccade is small, and less or not at all when the distance is large.

For each of the 4 monkeys, the effect of the speed of the new target motion on post-saccadic pursuit depended on the distance from the endpoint of the saccade to the target. For the smallest distances of about 1 degree, the post-saccadic eye velocity was highest for new target motion at 30 deg/s (Figure 1.5B, gray symbols), intermediate for new target motion at 20 deg/s (white symbols), and lowest for new target motion at 10 deg/s (black symbols). For the largest distances of about 9 degrees, the post-saccadic eye velocities are much more similar for the three target speeds. For the intermediate distances, the separation of the responses for the three target speeds also was intermediate. We performed two statistical analyses to quantify our observations. First, we compared the y-intercepts of regression lines fitted to the relationship between post-saccadic eye velocity and distance from saccade endpoint to target for each of the three final target speeds. When the pre-saccadic target speed was 10, 20, or 30 deg/s, the y-intercept was 9.05 ± 0.85 , 13.45 ± 0.67 ,




Figure 1.5: Effect of the distance from the endpoint of the saccade to the target position on the relationship between post-saccadic eye velocity and target speed. A: Schematic diagram showing the task as a sequence of 4 snapshots in time. A1. A target (open circle) began moving at the fixation point at 20 deg/s. 2. The target stepped and moved at a speed indicated by the marker: 10 deg/s (black filled circle), 20 deg/s (open circle), or 30 deg/s (gray circle). Only one target was visible to the monkey on the trial. The eye (x) begins pursuit in the direction of initial target motion. A3. The eye makes a saccade to the old location of the target and moves in the same direction. Only the 20 deg/s target is shown for clarity. A4. The eye makes a saccade to the new location of the target and tracks it appropriately. (B) Four graphs plotting the mean and standard deviation of the absolute value of post-saccadic eye velocity against the distance between target and eye, for three different speeds of target motion. Different axes and markers correspond to different monkeys. Black filled markers are trials where the target slowed to 10 deg/s, open markers where the target continued to move at 20 deg/s and grey markers where the targets sped up to 30 deg/s.

and 20 ± 3.56 . All differences in y-intercept were statistically significant (t-test, $p < 0.05$). Second, we ran an ANOVA using the speed of new motion (d.f = 2) and the size of the target step (d.f. = 2) as factors on all data points ($n = 2928$). There was a main effect of the size of the target step, and no main effect of speed ($F = 6.29$, $F = 0.53$; $p < 0.01$, $p = 0.53$). Supporting our observations, there was an interaction between the size and speed ($p < 0.01$). Post-hoc comparisons across factors failed to reach significance.

The effect of target speed on the post-saccadic eye velocity for a given distance between saccade endpoint and target is diagnostic for how strongly the new visual motion just before the saccade affects pursuit. Thus, the results in Figure 1.5 indicate that speed is processed within the spatial filter for pursuit, as is direction.

1.4.6 Frame of reference of the spatial filter for pursuit.

The experimental design used so far shows that there is a spatial filter for pursuit, as the post-saccadic eye velocity depends on the location of the moving target relative to the endpoint of the saccade. However, it does not resolve the frame of reference of the spatial filter. In all of our experiments so far, the target was displaced away from both the endpoint of the saccade and the current location of the fovea. As such, our previous experiments cannot disambiguate whether the spatial filter is positioned in relation to the location of the target on the retina before or after the saccade. The experiment diagrammed in Figure 1.6A can do so.

Each trial began with the monkey fixating straight ahead. At the start of interval #1, the fixation target disappeared and another target appeared eccentrically at 3 degrees up and right. At the start of interval #2, before the saccade, the target stepped 3 degrees up or down and began to move back toward the vertical meridian either along the horizontal axis (Figure 1.6A, interval #2b, "along") or along a parallel line 6 degrees off the axis (Figure 1.6A, interval #2a, "above"). Once the monkey makes a saccade to the original location of the target at 3 degrees up and right, these two options disambiguate the coordinate system

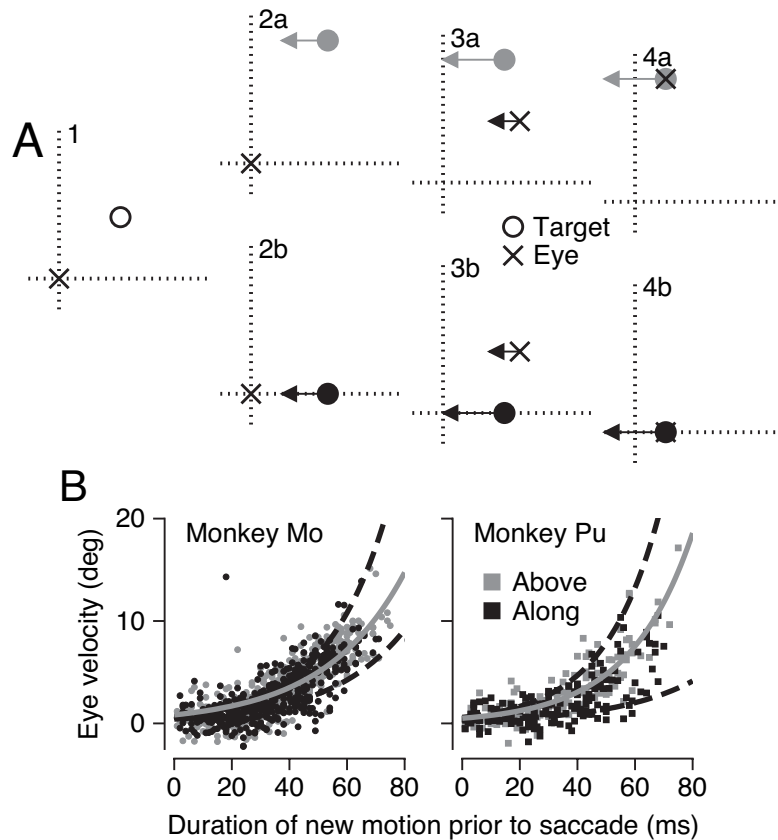


Figure 1.6: Analysis of the frame of reference of the relationship between post-saccadic eye velocity and the distance from the endpoint of the saccade to the target position. A: Schematic diagram showing the task as a sequence of 4 snapshots in time. A1. A stationary target (open circle) appeared 3 deg above and to the right of the eye (x). A2a/b. The target could follow onw of two trajectories. It could step 3 deg upwards and begin moving to the left, "above" trials, indicated with a grey filled circle. Or, it could step 3 deg downward to the horizontal meridian and begin moving to the left, "along" trials indicated by black filled circles. A3a/b. The eye makes a saccade to the original location of the target, and begins to move smoothly to the left. A4a/b. The eye makes a saccade to the correct location of the target and continues to track it smoothly. (B) Graphs that plot eye velocity as a function of the duration of target motion prior to the first saccade for two monkeys. Grey and black dots show data from individual trials in which the target stepped so that it moved either above or along the horizontal meridian. The solid grey curves show the exponential fit to the "above" trials, while the two dashed black curves delimit the 95% confidence intervals on the fit to the "along" data.

of the spatial filter. For both the "above" and "along" target motions, post-saccadic eye velocity is driven by new target motion that is approximately 3 deg from the eye position at the endpoint of the saccade. If the spatial filter is centered on the endpoint of the saccade, then the responses to either target motion should be the same, because in both cases, the motion is 3 degrees from the endpoint of the saccade. However, the "above" and "along" target motions provide new target motion at quite different retinal eccentricities just before the saccade. The target motion along the horizontal meridian is very close to the fovea while the target motion that is above the meridian is more than 6 deg eccentric. If the filter is centered in retinal coordinates on the fovea just before the saccade, then post-saccadic eye velocity should be larger for the target that stepped back to the horizontal meridian.

Post-saccadic eye velocity was the same for both target trajectories. The symbols in Figure 1.6B plot two monkeys' post-saccadic eye velocity on individual trials, as a function of the duration of the new target motion that caused the pursuit response. The duration was defined as the interval between the step and ramp of target position at the start of interval #2 and the monkey's saccade at the start of interval #3. Post-saccadic eye velocity increased as a function of the duration of motion exposure, and the data overlapped extensively for the along target motion (black symbols) and above target motion (grey symbols) for both monkeys. To quantify this claim, we fitted an exponential to the data (1.2), and calculated the 95% confidence intervals for our fit. The grey line representing the fit to the "above" trials sits well within the black dashed lines, which show the 95% confidence intervals of the fit to the "along" trials, demonstrating that the response of the pursuit system under these different conditions cannot be distinguished. For monkey Mo, the latency (β) and time constant (τ) were 0.19 and 0.04 for 471 "along" trials; for 436 "above" trials, β and τ were -0.64 and 0.04. For monkey Pu, β and τ for 148 "along" trials were -0.69 and 0.05; for 153 "above" trials, β and τ were -0.99 and 0.04. We conclude that the spatial filter for pursuit is centered on the endpoint of the saccade: the location of new motion relative to the endpoint of the saccade is more important in determining the magnitude of pursuit

than is the actual retinal location of the target.

1.4.7 Dynamics of processing in the spatial filter for pursuit

Figure 1.6B shows that post-saccadic eye velocity depends strongly on the duration of the new motion before the saccade. Figures 1.7A-F analyze the dynamics further by re-plotting the data from the experiments reported in Figure 2 and the accompanying text. Recall that in these experiments, the targets were displaced and assumed an orthogonal direction of motion at the start of interval 2. For each of the three different step sizes (0/1, 5, or 9 deg) in that experiment, we plotted horizontal eye velocity when it was in the direction of new target motion in Figures 1.7A-C, and when it was in the direction of old target motion in Figures 1.7D-F. The gray clouds of points show normalized post-saccadic eye velocity as a function of the duration of new motion for each individual trial from 4 monkeys. Eye velocity was normalized relative to the average post-saccadic eye velocity measured on control trials in the same experiment, when the target started to move horizontally and continued without any changes.

Figures 1.7A-C show that longer durations of pre-saccadic target motion in the new direction caused larger post-saccadic eye velocities. The duration of the rise from zero to 90% of maximal response is between 90 and 130ms, and is similar to the initial response of the smooth pursuit system to optimal step-ramp target trajectories. For a direct comparison between the two, see Figure 5 of [Lisberger, 1998]. The data are well-approximated by the sigmoid curve in each graph; the solid ribbon shows 95% confidence intervals around the best fitting curve. Details about the fits and the exact values of the sigmoid at 90% can be found in Table 2. Figures 1.7D-F show that post-saccadic eye velocity in the direction of original motion declines as a function of the duration of the new motion. The data could again be summarized by the superimposed decreasing sigmoid functions plotted as ribbons to show their 95% confidence intervals. The functions describing the decrease in post-saccadic eye velocity are related to the offset of the pursuit response, as analyzed in




Figure 1.7: Temporal dynamics of the modulation of visual-motor transmission for pursuit. All plots show normalized horizontal post-saccadic eye velocity as a function of the duration of new motion before the beginning to the targeting saccade. Normalization was performed relative to the post-saccadic eye velocity on control trials in which a single target moved horizontally. Graphs are grouped by rows, which show a particular type of experiment, and columns, which show data when the displacement of the target was 1, 5, or 9 deg. In each graph, the tiny gray dots show the responses in individual trials across all 4 monkeys. A-C: Data from the experiment described in Figure 1-2. Initial target motion was vertical and the new target motion was horizontal. The graphs plot the time course of the increase in post-saccadic eye velocity in the direction of new motion. D-F: The original direction of motion was horizontal and the new direction of target motion vertical. The graphs plot the time course of decay of the response to the old, horizontal target motion. In A-F, the black ribbons show the best fitting sigmoid function with the 95% confidence intervals on the fit. G-J: Data from the experiment diagrammed by open circles in Figure 5A, where the target underwent a vertical displacement but continued to move horizontally at 20 deg/s so that the original and new motions were the same. The jagged solid line corresponds to the mean of the data smoothed by a 5 ms sliding window. The dashed line shows to the sum of the sigmoid fits from the data in the in two graphs directly above each of the bottom graphs.

Figure 4 of [Luebke and Robinson, 1988], Figure 5B of [Soechting et al., 2005], and throughout [Krauzlis and Lisberger, 1994]. The parameters of the sigmoid fits to the data (1.1) reveal how the location of the visual stimulus within the spatial filter affects the dynamics of the gain of visual-motor transformations for pursuit. As shown in Table 1.2, the slope decreased and the half-maximum time of the best fitting sigmoid increased as a function of the distance from the saccade endpoint to the target. Thus, gain of visual-motor transformations rises more quickly to a higher level at the center of the spatial filter.

The time courses of gain modulation in Figure 1.7 allow us to explain why the values of post-saccadic eye velocity in the direction of the original target motion were small and not strongly related to the duration of new motion in Figures 1.2D and 1.3D. Figures 1.7D-F indicate that post-saccadic eye velocity in the original direction of target motion decays quite rapidly, reaching 10% of its maximum value when there had been as little as 65 ms of new motion (Table 1.2). Because we included data in our graphs only if there had been in excess of 50 ms of new motion, the values of post-saccadic eye velocity should be small and therefore the effects of the distance from the endpoint of the saccade to the target difficult to see.

Table 1.2:

Step size	Slope (95% CI)	Half-max (95% CI)	# Trials	Time at 0.10	Time at 0.90
0/1 deg, "New"	0.061 (0.003)	53.9 (0.8)	2814	18.94	90.01
5 deg, "New"	0.053 (0.003)	73.4 (1.36)	3055	30.67	112.82
9 deg, "New"	0.045 (0.003)	86.3 (2.23)	3161	37.04	129.54
0/1 deg, "Old"	-0.066 (0.003)	36.0 (0.7)	3446	64.24	7.39
5 deg, "Old"	-0.050 (0.002)	52.4 (0.88)	4737	88.89	9.03
9 deg, "Old"	-0.048 (0.003)	58.9 (1.21)	4234	97.38	12.76

1.4.8 Motion processing within the spatial filter reflects the weighted vector superposition of old and new target trajectories

The bottom row of graphs in Figure 1.7 plots data for the trials from Figure 1.5, where the target moved horizontally at 20 deg/s, and underwent a step displacement but did not

change direction or speed. Even though the same direction and speed of target motion was present both before and after the step change in target position, the clouds of points in Figures 1.7G-I suggest a parabolic shape that is clarified by averaging the normalized eye velocity in 1 ms bins to yield the functions shown by the bold, jagged curves.

The transient dip in the relationships between post-saccadic eye velocity and the duration of the visual motion stimulus even when the target has not changed speed or direction suggests that a step of target position initiates two simultaneous components of the pursuit response: decay of the response to the original motion and build-up of a response to the new target motion. To test this hypothesis, we added the paired functions derived from the data in the first two rows of Figure 1.7, which provide direct estimates of the time courses of build-up and decay of the horizontal pursuit response. For each size of target step, the resulting curves (Figure 1.7G-I, dashed curves) predict the means of the actual data quite well, diverging only as the number of points becomes small when the duration of "new" motion approaches 75-100 ms. The analysis in Figure 1.7 provides direct support for the conclusions of [Soechting et al., 2005] who showed that the pursuit response to an abrupt change in target direction could be described as the vector superposition of two discrete responses.

The graphs in Figure 1.7G-I explain why the average post-saccadic eye velocity did not depend on the distance from the saccade endpoint to the target when the target underwent a step of position but did not change speed or direction (Figures 1.5B-E, open symbols). Under these conditions, eye velocity was near its minimum when the duration of new target motion was 50 ms, which is the situation that obtained when many of the data in Figure 1.5 were measured. However, the minimum was not a large decrease from maximum, and did not depend on the distance from the endpoint of the saccade to the target position.

1.5 Discussion

We have demonstrated that target motion is processed preferentially to guide smooth pursuit eye movements within a symmetric spatial region around the endpoint of a saccade. When the initial target motion at one location in one direction is followed closely by new target motion at a new location, the resulting pursuit eye movement reflects a weighted vector superposition of the two. The weighting is related to, and varies dynamically in time as a function of, the spatial distance between the saccade's endpoint and the target. Intuitively, visual motion signals stimulating the foveal retina ought be more effective in driving pursuit; given the latency of visual processing, this effect would first be detectable 100 ms after the saccade brings the eye to point at the target. Instead, our results demonstrate that the weighting change is present immediately after the saccade and therefore, because of the latencies in sensory-motor processing for pursuit, must begin before the saccade. Thus, the pursuit system anticipates the spatial location of target motion with respect to the endpoint of a saccade, and responds unhampered by the visual latency. Our Discussion will make the case that this very specific oculomotor interaction affords a comprehensive and mechanistic description of gain control for smooth pursuit eye movements.

1.5.1 Old and new concepts of gain control for pursuit

There is a marked difference between pre-saccadic and post-saccadic smooth eye velocity when comparing responses at comparable times after the onset of target motion [Lisberger, 1998, Ogawa and Fujita, 1998]. Lisberger [1998] suggested that post-saccadic enhancement could be thought of as one consequence of a "volume control" turned to a high setting by the execution of a saccade. Gardner and Lisberger [2001, 2002] demonstrated that the saccade-driven increase in pursuit gain was inextricable from more cognitive processes of decision-making. When confronted with the motion of two targets, pursuit generates an initial, pre-saccadic eye velocity that is a weighted vector average of the independent

responses to the two targets. A natural or electrically evoked saccade to one target is sufficient to drive post-saccadic smooth pursuit preferentially by the motion of the target closest to the endpoint of the saccade. The selective change in pursuit gain is sufficient to instantiate a "choice" between the two targets.

We propose a single model of pursuit that can explain our data as well as a two prior observations: smooth pursuit results from a stereotyped response to visual motion, where the response is modulated by a spatial filter that operates as though it moves with the eye. First, Lisberger and Westbrook [1985] demonstrated the existence of a comparable spatial filter in their exploration of the consequences of initial target eccentricity on pre-saccadic pursuit. They attributed the modulation to weaker visual motion signals from more eccentric retinal locations. Even though the filter described in the present paper is centered on the endpoint of the saccade, rather than the current position of fixation, their filter is remarkably similar. Second, the temporal dynamics of the filter, revealed in our Figure 1.7 and Figure 5 of Lisberger [1998], follow a time course remarkably similar to that of the "best" pre-saccadic eye velocity initiated by a target that ramps toward the position of fixation from just eccentric [Lisberger and Westbrook, 1985]. Given the similarity of the spatial and temporal properties of the filters that affect both pre-saccadic and post-saccadic pursuit, it seems reasonable to conclude that they are mediated by the same neural mechanisms. The saccade plan represents the upcoming locus of the fovea; as the plan is formulated and executed, the locus of the strongest positive modulation of visual motion inputs moves to the upcoming location of the fovea.

Liston and Krauzlis [2003, 2005] have proposed a different explanation based on data that fit well with ours. They used the shorter latency of pursuit to study the response of the oculomotor system to two targets. They contrived a situation where the initial pursuit response could be in a different direction than the eye movement following a saccade, and studied pursuit dynamics in the pre-saccadic period. They posit that a single high level "decision process" coordinates the tracking response, forcing the smooth eye move-

ment to match that of the target closest to the endpoint of the saccade. Their data are consistent with our proposal for the role of the saccade: the formulation and execution of the saccade plan itself provides the "decision signal" for the pursuit system by enhancing the transmission of visual motion signals from the region of space at the endpoint of the saccade. Further, the time-course of their effect resembles the formulation of a plan to saccade. Our interpretation is consistent with their report that subthreshold microstimulation in the superior colliculus can bias pursuit toward the motion of targets passing through the movement field of the excited neurons [Carello and Krauzlis, 2004]. Finally, our interpretation provides a mechanistic explanation for why space ought be a primal cue for the pursuit system, as Adler et al. [2002] have described: a spatial cue provides direct information to the pursuit system about the upcoming location of the fovea, just as does the saccade plan. We conclude that our data provide a quantitative description of the influence of the upcoming locus of the fovea, as driven by the saccade plan, on smooth pursuit eye movements.

1.5.2 Features of the neural underpinnings of a moveable spatial filter, and perceptual consequences

We propose that weighting retino-centric visual motion by the upcoming position of the eye in space transforms the representation of visual motion into a spatial command for eye velocity. Saccades rely on a similar coordinate transformation, taking account of the former retinal position of a flashed stimulus and all subsequent changes in the position of the eyes [Mays and Sparks, 1980]. Insights from computational studies provide a framework for thinking about how this kind of computation might be done in the brain. If model neurons combine extra-retinal information about the locus of gaze with visual signals about the retinal position of a target, then they can learn and compute a movement command appropriate for bringing the eyes to the position of a target in space [Zipser and Andersen, 1988, Pouget and Sejnowski, 1994, Salinas and Abbott, 2001, Smith and Crawford, 2005].

These models suggest that neurons involved in creating a spatial-weighting of the visual motion drive for pursuit ought integrate two signals: A directional signal related to the retinal visual motion stimulus, and a direction-independent gain modulation related to eye position. The effects of electrical stimulation of pursuit-related neurons in the frontal eye fields were interpreted as having two similar components, suggesting a direct neural correlate of the interactions needed to explain our data [Tanaka and Lisberger, 2002a].

It would be tempting, but unproductive, to speculate about the anatomical loci of either the filter itself or the signals that control it. It is unlikely that neural correlates of such fundamental sensorimotor transformations would be restricted to a single anatomical locus, as neurons across the sensorimotor spectrum show spatially-selective enhancement in conjunction with saccades [Newsome et al., 1988, Tolia et al., 2001, Reppas et al., 2002]. Even subthreshold micro-stimulation in an area that generates saccades (FEF) enhances visual responses in the extrastriate visual cortex (V4) to targets presented at the endpoint of saccades generated by stronger stimulation at the same site [Moore and Armstrong, 2003]. While it is likely that the filter's effects will be felt across many anatomical areas, research into the neural correlates of attention suggests that these effects might be different in magnitude in tractable ways.

Our proposed filter fits well with concepts suggested by prior research on attention. For example, neurophysiologists have demonstrated many correlates of perceptual enhancement where the neural responses to a given visual stimulus are larger if a monkey attends to the location of the stimulus or is going to use the stimulus as the target of an impending saccade, e.g. [Goldberg and Wurtz, 1972, Mountcastle et al., 1981, Connor et al., 1997, Bisley and Goldberg, 2006]). The magnitude of these effects can vary with the area. Consider two areas involved in representing retinal motion to drive smooth pursuit eye movements, MT [Newsome et al., 1985, Carey et al., 2005]) and VIP [Schlack et al., 2003]. The magnitude of neural modulation following cued attention is considerably larger in VIP than in MT [Cook and Maunsell, 2002]; crucially, VIP receives considerable informa-

tion about the position of the eyes, whereas MT has little or no extra-retinal input [Newsome et al., 1988, Bremmer et al., 1999, Cook and Maunsell, 2002]. Correspondingly, we predict that the neural correlates of the spatial control of visual motion processing for pursuit will be clearest in areas that have responses related to visual motion and/or pursuit, and also are driven strongly by information about the direction of gaze.

Lastly, we note an intriguing parallel between perceptual and motor effects of saccades. In the face of many saccades, we maintain a stable perception of the world. As Dodge [1907] observed, the inverse of the command to saccade is the exact signal needed to maintain such perceptual constancy. Perception of both space and time is compromised in specific ways around the time of the saccade [Morrone et al., 2005]. These perceptual effects have been modeled as resulting from a coordinate transformation similar to the one we proposed to explain our results [VanRullen, 2004]. Perhaps use of the pursuit system to obtain continuous measures of the neural dynamics surrounding saccades will allow an understanding at the level of neural circuits into this most fundamental and intriguing perceptual phenomena.

Chapter 2

Sparse Encoding of Smooth Eye Movement by Single Spikes in Frontal Eye Field Neurons

2.1 Abstract

Cortical neurons are customarily described as noisy encoders, whose responses must be pooled to precisely encode behavior. In this study, we show that spikes from single neurons in the pursuit sub-region of the macaque frontal eye field are sensitive enough to encode trial-to-trial fluctuations in smooth eye movements. Using linear filter estimation techniques, we model the co-variation of spikes and behavior, allowing us to measure and describe the nature and strength of the relationship. Individual neurons tend to have strong correlations at specific time periods, roughly in line with their moments of peak firing, which tile the duration of the trial in a sparse fashion. The shape of the filters suggests a network of orthogonal encoders, supported by the finding that filtration decreases the pairwise correlation between simultaneously recorded pairs. We estimate the co-variation of neural activity, and build a simple summative population model compatible with our

hypotheses of sparse network function. Our population model supports the idea that even in the face of co-variation between pairs of neurons, behavioral variability arises from the activity of a small set of highly active neurons.

2.2 Introduction

Cortical neurons respond to repeated stimulus presentations with highly variable patterns of spikes [Tolhurst et al., 1983], yet we are able to make incredibly sensitive perceptual judgments [Sakitt, 1972]. Some have proposed that this trial-to-trial variability represents noise in sensory encoding, and suggested that responses of many neurons must be pooled to obtain reliable perceptions [Shadlen and Newsome, 1998]. Recent theoretical work has suggested that variability in and among cortical neurons may represent a probabilistic estimate of how certain we are about a sensory stimulus [Ma et al., 2006]. Both these models suggest that single spikes play little role in the final perceptual judgment.

Like sensory neurons, cortical neurons in motor areas vary their firing from movement to movement. Unlike sensory stimuli that can be presented repeatedly, movements towards the same goal vary across repetitions. Does variation in the firing of motor neurons represent noise that is averaged out by pooling operations, or might this variation form part of the motor command signal itself? Recent work has successfully linked preparatory cortical dynamics to motor variability [Churchland et al., 2006a,b, Riehle and Requin, 1993], suggested a sensory origin for motor variability [Osborne et al., 2005], and begun to model the consequences of motor variability [Harris and Wolpert, 1998, Todorov and Jordan, 2002]. Work with large ensembles of simultaneously recorded neurons for continuous behavior suggests that neural variability in cortical motor areas is so potent that stable representation of movement is only possible by pooling large numbers of neurons [Carmena et al., 2005, Lee et al., 1998, Maynard et al., 1999, Paninski et al., 2004b]. To examine this suggestion, we explored whether single cortical spikes could encode the trial-to-trial vari-

ations present in a repeated oculomotor behavior.

Pursuit is a dynamic oculomotor behavior engaged when primates attempt to keep their fovea pointed at a target that is moving in space. A subset of neurons in the frontal eye fields (FEF_{SEM}) of the macaque monkey respond preferentially during smooth tracking in a given direction. Electrical microstimulation of the FEF_{SEM} both drives smooth eye movements and increases the gain of the eyes' response to target motion. Previous recording studies have demonstrated a link between the mean firing of single neurons in the FEF_{SEM} and ongoing smooth pursuit [Tanaka and Lisberger, 2002a]. We thus capitalize on the strong relationship between spikes in the FEF_{SEM} and the eye movement in the subspace of the neuron's preferred direction [Ringach et al., 1997].

To quantify the co-variation between neural activity and the dynamic pursuit behavior, we use linear systems analysis to estimate the nature and fidelity of the transformation between brain and behavior [Halliday et al., 1995, Paninski et al., 2004a]. We derive a function that represents the transformation between deviations from the "average" spiking activity of single neurons in the FEF_{SEM} and simultaneous deviations from the mean eye velocity, and show that this function can predict the trial-to-trial dynamics of the eye movement for a set of neurons. An individual neuron's peak activity is restricted to particular times in the trial, which we term "sparseness," and correlates well with the time of best correlation between neural and behavioral variation. We model the consequence of such a relationship, demonstrating the limits placed on the level and nature of downstream noise. Using measurements of pairwise covariation between filtered spike trains from simultaneously recorded neurons, we model the generation of an eye movement in a manner compatible with a sparsely driven population. Finally, classification of the filter shapes suggests an orthogonal basis set, and we demonstrate that the filters act to decorrelate the spike trains. Our data and models are compatible with the idea that the FEF_{SEM} contains a sparse representation of smooth eye movement, and that spikes from individual cortical neurons may play a strong role in driving the dynamics of ongoing behavior.

2.3 Results

Our experiments, analyses, and models were designed to quantify, at the level of short time intervals and single trials, the degree of co-variation between dynamic behavior and spikes from single and paired units in the macaque cortex. We begin with a description of the fluctuations present in both behavior and neural activity when two monkeys repeated a highly constrained oculomotor behavior. Next, we use linear filter estimation techniques, closely related to those used to characterize sensory responses, to model the relationship between neural and behavioral variation, and to quantify the precision with which neurons can encode variation. We describe the shapes and latencies of our population of filters. We next show the temporal heterogeneity of peak neural activity, and describe its relationship to the time course of the neuron-behavior correlation. To understand the conditions under which such a relationship obtain, we explore a simple model of the downstream noise. We then show data from a pair of simultaneously recorded units from our data set, and the relevant distributions we measure. A model of the population follows, using the distributions from all pairs to estimate the size and nature of the network. Finally, we conclude by demonstrating that the orthogonality of our filters serves to decorrelate spike trains.

2.3.1 Neural and behavioral variation expressed as residuals

To elicit a set of highly stereotyped behaviors, we trained two male monkeys to fixate on a target straight ahead, and then to track the target after rapid displacement to one side, followed by constant velocity in the opposite direction across the original position of fixation. This "step-ramp" stimulus naturally induces the eyes to track the target, beginning with a smooth acceleration of the eyes 100 ms after target motion begins, followed by sustained smooth tracking near the velocity of the target. Though the eye movement response to step-ramp target motion is stereotyped, animals show considerable variation in

the latency to move, the magnitude of acceleration, and the final tracking velocity (Figure 2.1A). The smooth component of eye movement evoked by the target motion has the same general form on each trial (Figure 2.1A, gray traces), but the spread in the distribution demonstrates that each motor response is unique.

Likewise, single neuron activity recorded in the FEF_{SEM} during repeated responses to the same target motion shows a stereotyped profile with subtle trial-to-trial variability. The raster in Figure 2.1B shows neural activity associated with each of the 95 eye movements included in Figure 2.1A. On every trial, the neuron was relatively silent before the eyes began to move, and was then generally active during tracking. For this particular neuron, activity was most vigorous in two intervals, one at the beginning of the movement, and one towards the end of the illustrated interval. Even in the face of this consistency, there is considerable variation in the precise number and timing of spikes on individual trials. The variation appears to be random in the sense that the raster does not suggest any consistent drift in the neural response across the duration of the experiment (from bottom to top of the raster). Consequently, we are confident that we can consider the response to be stationary across the experiment.

To examine behavioral and neural variation in more detail, we plotted the residuals of eye movement and firing rate, defined as the difference between the response on each individual trial and the mean across trials. The residuals of eye velocity (Figure 2.1C, gray traces) reach maxima of about +5 deg/s, about twice the magnitude of the standard deviation of the original traces (black trace) and considerably smaller than the 15-20 deg/s of eye velocity in the original traces. Inspection of Figure 2.1C, and more detailed quantitative analysis, confirm that the residuals form a continuous distribution that is fairly uniform across the part of the trial that contains smooth tracking. While some of the variation no doubt arises from the way in which we've aligned the trials, all of our findings are robust to trial alignment. Similarly, our findings obtain when we perform our analysis on the eye position trace, rather than on eye velocity records (Figure 3.1).

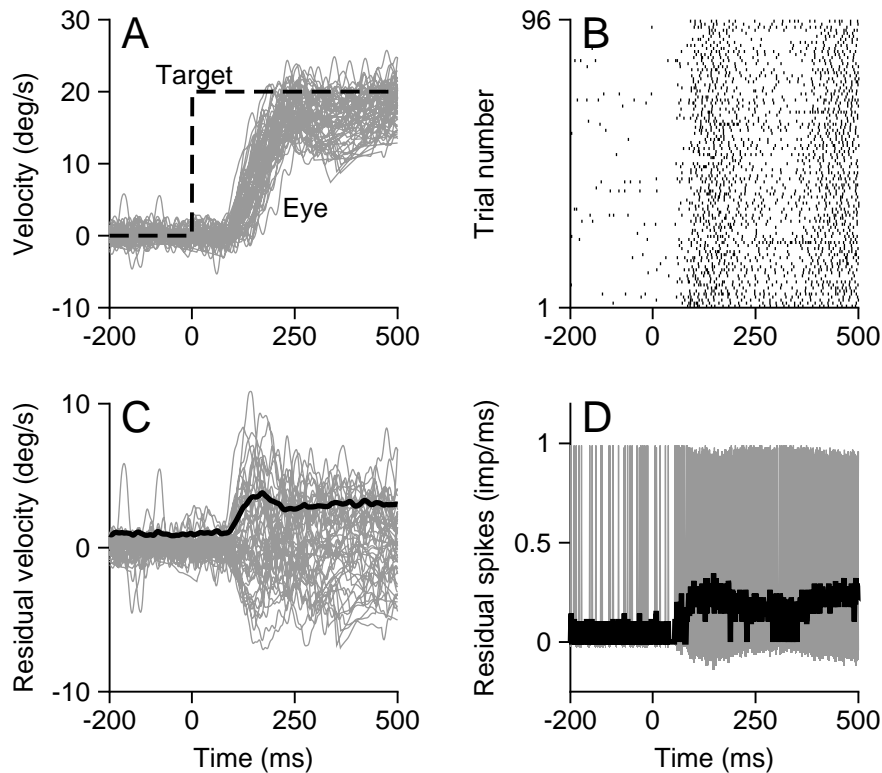


Figure 2.1: A typical example of neural and behavioral variability. A, The gray area shows superimposed eye movements in response to 95 presentations of identical target motion. The dashed black line indicates target velocity as a function of time. B, Raster showing the responses of a single neuron during the eye movements shown in A; each black tick indicates one spike. C, Residuals of eye velocity derived from the data in A by subtracting the mean of all responses to a given target motion from each individual trace. The black trace shows standard deviation of the residuals as a function of time. D, The gray area shows the residual spikes as a function of time; the thick black line shows the standard deviation of the firing rate residuals.

We obtained an analogous representation of the variation in firing rate by starting with a binary representation of each trial. In each 1 ms-wide bin, the neural activity takes on a value of one, if a spike occurred in that bin, or zero, indicating no activity. To compute the residuals of firing rate, we subtracted the average activity across trials (effectively, a per-stimulus time histogram) from the binary representations of the neural activity on each trial. When the residuals from all 95 trials are superimposed (Figure 2.1D), the spikes from individual trials cannot be resolved but it is possible to see important trends. Within each individual trial, there is a small but non-zero probability in any 1-ms bin that a spike will have occurred. When a spike is present, it takes the gray traces up to a value of one minus the mean across trials. Because 95 traces have been superimposed, the spikes appear as a positive gray area during the time when the neuron responds. In most bins, however, a spike will not have occurred so that the value of the residual is the negative of the mean firing rate ($\langle fr \rangle$), or the PSTH, inverted. At higher resolution (not shown), it becomes clear that the top and bottom of the gray region ride, as expected, along $(1 - \langle fr \rangle)$ and $(-\langle fr \rangle)$. This and other analyses were done with a bin size of 1 ms. We show in Supplementary Figure 3.2 that altering the bin size, or sampling resolution, does not change the nature of our results.

Figure 2.1 illustrates that there is a reasonable degree of trial-by-trial variation in both smooth eye movements and neural activity when the monkey tracks the same target motion repeatedly. The heterogeneity provides a substrate for asking whether there is a relationship at the level of individual trials between neural and behavioral fluctuations.

2.3.2 The structure of correlations between neural and behavioral variation

To ask whether eye movements co-vary with single neuron activity, we must first derive a relationship between spiking and movement. Spikes in FEF_{SEM} could be correlated with changes in the eye movement that precede or follow them, and might be related to any number of movement parameters such as velocity or acceleration. We adopt a linear sys-

tems approach that models the eye movement velocity residuals on each trial as a linear filter of the spiking residuals. This approach makes few assumptions about what the neuron encodes, and can capture relationships between spiking and any linear combination of movement parameters. Similar to filters derived from reverse correlation in sensory systems, the filter for any particular neuron provides a linear approximation of "what the neuron encodes." Using residuals restricts the analysis to a small range of eye movements, with the advantage of making our linear model a reasonable approximation to the complete non-linear relationship between spiking and movement.

The linear filter is closely related to the more familiar description of the trial-by-trial correlation between the residual firing rate and the residual eye velocity (Brenner et al., 2000; Gerstein and Perkel, 1972; Vaadia et al., 1995). In the correlation "map" shown in Figure 2.2A, each pixel represents the strength of the relationship between one time point in the eye movement and one time point in the neural response. The color of each pixel represents the value of Pearson's correlation coefficient between the residual of eye velocity and the residual firing rate across all trials, evaluated at all pairs of time points. To make the trends in the data more readily apparent, we've smoothed the residual spike train with a Gaussian function of width 20 ms.

The upper-left corner of Figure 2.2A represents the period of fixation before the target starts to move (c.f. Figure 2.1A), and is mottled green and blue, corresponding to the limited variation in both the eyes' velocity and the neural response. During the initiation of pursuit ($t = 100$ to 200 ms), there is a strong positive correlation between the residual firing rate, and the residual eye movement 10-20 ms later. The correlation appears as the red peak with a dark core just below the black diagonal line. The correlation map also contains a wide region of positive correlations at later times in the response, and a blue strip indicating strong negative correlation just to the right of the strongest positive correlation.

The hotspots in the correlation map of Figure 2.2A suggest that the neuron's response across trials does co-vary with the animal's behavior. However, the analysis has sev-

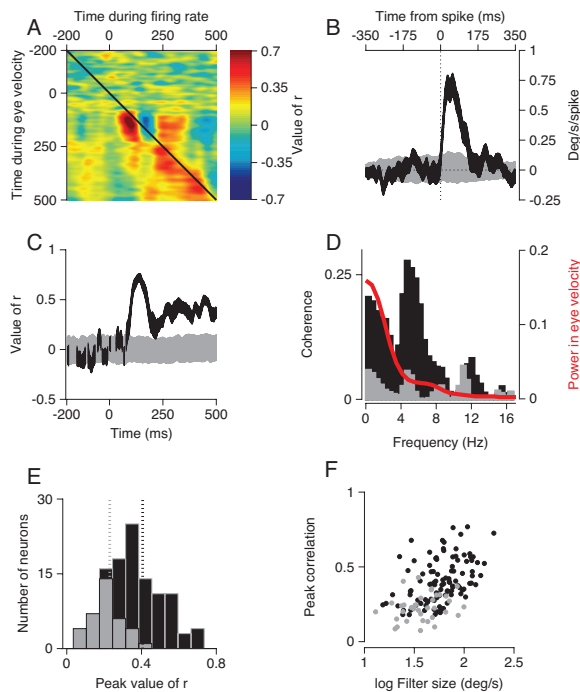


Figure 2.2: The co-variation between spikes and eye movements for an example neuron. A, The joint peri-behavior correlation histogram between the eye movements and spikes. Each pixel shows the correlation coefficient between the eye velocity and the spikes at times given by the position of the x- and y-axes. B, The black trace shows the filter that describes the relationship between residual eye velocity and spikes as a function the time elapsed from a spike. Negative and positive values of time correspond to changes in eye velocity that precede or follow the spike. The gray trace shows the corresponding filter derived from shuffled trials. The width of the function represents one standard deviation above and below the distribution of filters derived from 150 draws of 40% of the trials. C, The black and gray traces plot the correlation coefficient between the actual and predicted residuals as a function of time, for the filters obtained from the actual and shuffled data, respectively. The width of the line represents one standard deviation above and below the mean for 150 tests of the 60% of the data reserved to test the filter. D, The coherence of the filter in B is plotted as a function of frequency; black corresponds to the normal transfer function, and gray to the shuffled. Overlaid in red is the normalized power in the residual eye velocity. E, Peak neuron-behavior correlation for each cell in the population; significant cells ($n = 105/141$) in black, non-significant cells in grey. Means for each distribution are plotted as lines, colored accordingly F, Filter size, on a log scale, plotted against the peak values for neuron-behavior correlation. Significant cells in black, non-significant cells in grey.

eral limitations: it does not provide a single number that encapsulates the amount of co-variation; it does not correct for the influence of correlation within the eye movement or spike train; and it does not provide a testable model of the relationship between spiking and behavior.

The filter shown in Figure 2.2B (black ribbon) corrects the limitations inherent in the correlation map of Figure 2.2A. It is derived by cross-correlating the residuals in the eye's velocity with the residuals in the spike train (Section 2.5), and is closely related to the average value along the diagonals of Figure 2.2A. The filter shows the time course of the variation in eye velocity that is related to the occurrence of one extra spike above the mean in our example neuron. The shape of the filter suggests that following each "extra" spike, the eye increases its velocity over a duration of ~ 130 ms, with a peak response of ~ 0.75 deg/s. The time where the filter peaks, 57 ms after the extra the spike, represents the time, probabilistically speaking, when the influence of an "extra" spike will be felt most strongly. The filter was corrected for the first-order correlations in the spike train; correcting for correlations in the eye movement yielded a similar filter, reflected about the zero-lag point. The gray ribbon describes a filter derived similarly on the same data, but shuffled to remove the trial-to-trial correspondence.

The filter in Figure 2.2B allows us to assess the magnitude and significance of neuron-behavior covariance, and provides a model of how an additional spike in this cell influences the behavior, but it does not tell us what percent of the behavioral variation can be accounted for by this neuron's firing. To obtain such a measure, we used the filter to predict the time course of the residual eye velocity for each test trial and computed the trial-by-trial correlation between actual and predicted residual eye velocity as a function of time (Figure 2.2C, black ribbon). Here, the black ribbon is visible only when the correlations were statistically different (paired t-test, $p < 10^{-3}$) from the analysis based on shuffled data (gray ribbon). The correlations were small and mostly non-significant before the onset of pursuit. They rose rapidly around the initiation of pursuit, reached values as

high as 0.7 for an interval of 20 ms duration, and were greater than 0.6 for an interval of over 50 ms. For the example neuron, maximal correlations between actual and predicted residuals of eye velocity fell around the time when the neuron fired most vigorously (cf. Figure 2.1B).

While corrected for correlations in the spike train residuals, our filter is still compromised by the correlations in the eye movement. To be sure of the veracity of our estimates of the relationship between neuron and behavior, we computed the coherence of our filter, plotted in Figure 2.2D as black bars. The coherence between two signals describes the variance at a given frequency accounted for by the function relating the signals (our filter); it is weighted by the autocorrelation of both eye movement and spike train residuals. If the coherence were one, there would be a perfect trial-by-trial correspondence between the neural and behavioral variation; the coherence is thus similar to the square of the correlation coefficient.

For the example unit, over the frequency range that contained most of the power in the eye movement residuals (0-8 Hz), the coherence averaged 0.25; compare to the coherence of the shuffled trials, plotted in grey. The coherence has two major peaks: one peak is at frequencies below 4 Hz where eye movement has the most power (red curve); the other peak is at ~ 6 Hz and probably corresponds to the initiation of pursuit, expressed as a small deflection at ~ 6 Hz in the power spectrum of the variation in eye velocity (red curve). The spectrogram of the eye movement residuals confirms this hypothesis (data not shown). The difference in the relative size of these two peaks in the eye movement and the coherence suggests that the neuron responds preferentially to more rapid changes in the eye movement, such as pursuit, and responds less to slow changes in the eye movement. Thus, the coherence agrees both quantitatively with the filter analysis, in that the filters account for about 25% of the variation in the actual residual eye velocities, as well as qualitatively, as the preferred frequency range matches the correlation obtained in the time domain analysis across the interval of pursuit (100 to 500 ms) (Figure 2.2C).

Population data confirmed the patterns seen in our example cell. Of 141 neurons in our sample, 105 yielded filters that were statistically different from filters based on shuffled data. For each neuron, we evaluated the correlation between the actual and predicted residuals of eye velocity as a function of time, as in Figure 2.2C, and measured the peak value of the correlation coefficient. For the neurons with significant filters (Figure 2.2E, black versus gray bars), the peak value of the correlation coefficient averaged 0.41 (range 0.18 to 0.76, SD: 0.15) versus those that were not statistically significant, averaging 0.23 (range 0.07 to 0.43, SD: 0.08). The correlations between actual and predicted residuals were significantly different for significant versus non-significant filters (t-test, $p < 10^{-9}$). The measures of coherence over the frequency range from 0 to 10 Hz (data not shown) confirmed those in the time domain: the mean for significant cells was 0.15 (SD: 0.10), while for non-significant cells the mean was 0.11 (SD: 0.08) (Wilcoxon rank-sum, $p < 0.02$). As expected, many neurons had values of coherence that peaked between 0 and 10 Hz (data not shown), in good agreement with the range of frequencies found in the eye movements.

Our population of neurons is not equivalent in their ability to encode behavioral variability. Figure 2.2F plots the peak neuron-behavior correlation against the log of the integral of the absolute value of the filter (a measure of filter "size"), with significant cells in plotted in black, and non-significant in grey. The log of the size of the significant filters is normally distributed (Bera-Jarque test, $p = 0.4$). The correlation coefficient between the peak neuron-behavior correlation and the size of the filter is significant, confirming that the larger filters are more reliable encoders of the trial-to-trial fluctuations in behavior ($r = 0.43$, $p < 10^{-5}$). Crucially, the logarithmic distribution shows that some neurons are considerably larger than others, allowing them to directly encode a greater range of dynamic fluctuations in the behavior.

We have illustrated that we can describe comprehensively, on a trial-by-trial basis, the nature of the correlation between the activity of a single neuron and the associated behavior. We extend the analysis to a population of neurons, and find that many can encode a

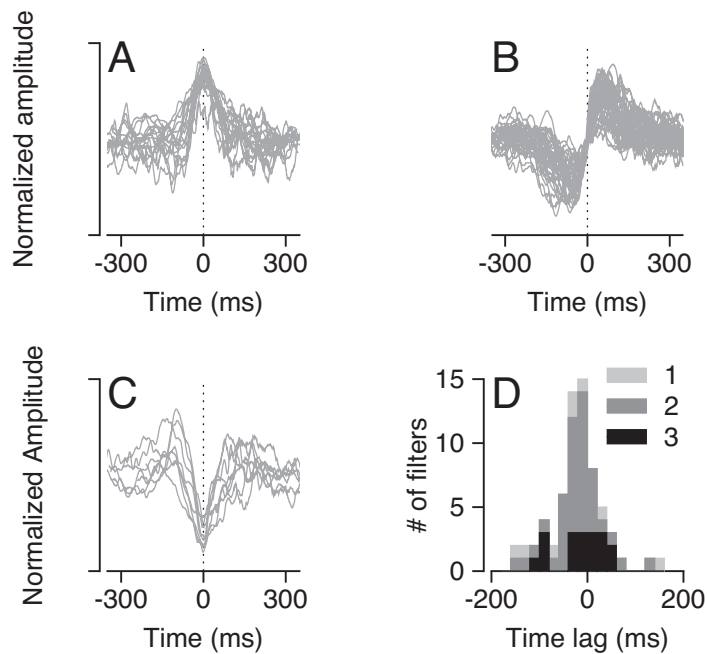


Figure 2.3: Properties of the filters that link neural activity to behavior. A,B,C, Superimposed time courses of statistically-significant filters that could be classified as monophasic (A, $n = 18$), biphasic (B, $n = 40$), triphasic (C, $n = 8$), or unclassifiable ($n = 39$). For clarity, filters are aligned at their center. D, The distribution of latencies is plotted as a stacked bar histogram for the three filter types shown in A-C. Monophasic filters are plotted in light gray, biphasic in dark gray, and triphasic in black.

surprisingly large fraction of the variability in the eye movement.

2.3.3 Filter shapes suggest an orthogonal basis set for movement commands

Do our neurons encode similar aspects of behavioral variation? The shape of the filter can be thought of as the "feature" that best describes the relationship between neural activity and the attendant behavior. As such, filter shape allows considerable insight into the behavioral consequences of a spike in a given neuron.

Of the 105 filters that were statistically significant, 66 could be classified according to their shape into three classes: 18 had a single, positive-going lobe (Figure 2.3A), 40 were biphasic (Figure 2.3B), and 8 were triphasic, with a large negative deflection (Figure 2.3C).

Each filter shape suggests a different role in encoding smooth eye movement. A single lobed filter describes an integrator: a neuron whose activity is related to a pulse of eye velocity. The width of the filter reflects the downstream smoothing of the impulse provided by a single spike in the neuron. A biphasic positive-going filter describes a situation where spikes are associated with a transition from lower to higher eye velocity, i.e. eye acceleration. Finally, a triphasic filter with a downward deflection implies that a spike is associated with an increase in the rate of change of the acceleration, or "jerk."

The latency and width of the filter provide another way to classify the data. Filters with a negative, or acausal latency are traditionally described as "sensory," inasmuch as their response is entirely determined by the information available before the spike. Filters with a positive latency are causal, and traditionally called "motor," as their effects are apparent only after the moment of a spike. The filter width suggests the time over which the relationship between neuron and behavior is present. We estimate the latency of each filter as the time from the extra spike to the "moment of maximal influence", defined as the time of the peak of an single or tri-lobed filter and the zero crossing for a biphasic filter. The distribution of time lags from the occurrence of an extra spike to the moment of maximal influence had a single peak centered close to zero (Figure 2.3D). There was no difference in time lag among the three filter types, and the mean latency from the spike to the moment of maximum influence was -0.93 ms (SD: 54.9 ms). The filter width was on the order of 150 to 350 ms — the width is due to the structure of smooth eye movements, which by definition are highly correlated in time. There was no relationship between the shape of a filter and the time of its maximum influence.

To confirm our method of filter clustering, we aligned the classifiable filters on their latency, and decomposed the resulting matrix into its principal components. The decomposition confirmed our earlier classification, and is described in full in Supplementary Figure 3.6. Briefly, the shape of the each of the first three significant eigenvectors was either single, double, or triple-lobed, and the clusters that emerged when filters were projected into

this space precisely matched our classification. We thus feel confident that our grouping of filters by the number of lobes is a reasonable approximation of the space in which they operate.

2.3.4 The timecourse of neuron-behavior correlations suggests sparse encoding, with low downstream noise

We have described a heterogeneous population of neurons, both with respect to their strength and features. Are they similarly heterogeneous with respect to the timecourse of the trial? Peak neural activity was remarkably dissimilar across our population of neurons. Figure 2.4A plots the normalized PSTH from each cell as a color-coded line, ordered by the latency to peak. The large red stripe tracks the peak of firing across the duration of the eyes' movement, and both cells with significant (top section) and non-significant filters (bottom section) tile the space. The cells averaged 64.37 ms (SD: 40.93) spent at or above 80% of their peak rate. Similarly, Figure 2.4B plots the timecourse of the correlation between neuron and behavior, (cf. Figure 2.2C), for the population, again split into cells with significant (top section) and non-significant filters (bottom section) ordered by latency to peak correlation. Just as with peak neural activity, when the eyes are moving, a similar stripe is visible.

The similar tiling across the trial led us to ask whether the cells most active at a given time best predicted behavioral variability at that time. We found that this was true: the mean correlation between the two time-courses was $r = 0.40$, and the range of our estimates (SD: 0.35) gave us confidence that this was unlikely to be artifactual. As plotted in Figure 2.4C, the strength of the neuron-behavior correlation predicted the likelihood that said correlation would resemble the PSTH ($r = 0.54$, $p < 10^{-10}$). Therefore, there is a relationship between the peak activity and the peak correlation time, and it is strongest for the best-correlated neurons.

Cells in the FEF_{SEM} are many synapses away from the muscles, and so the transforma-

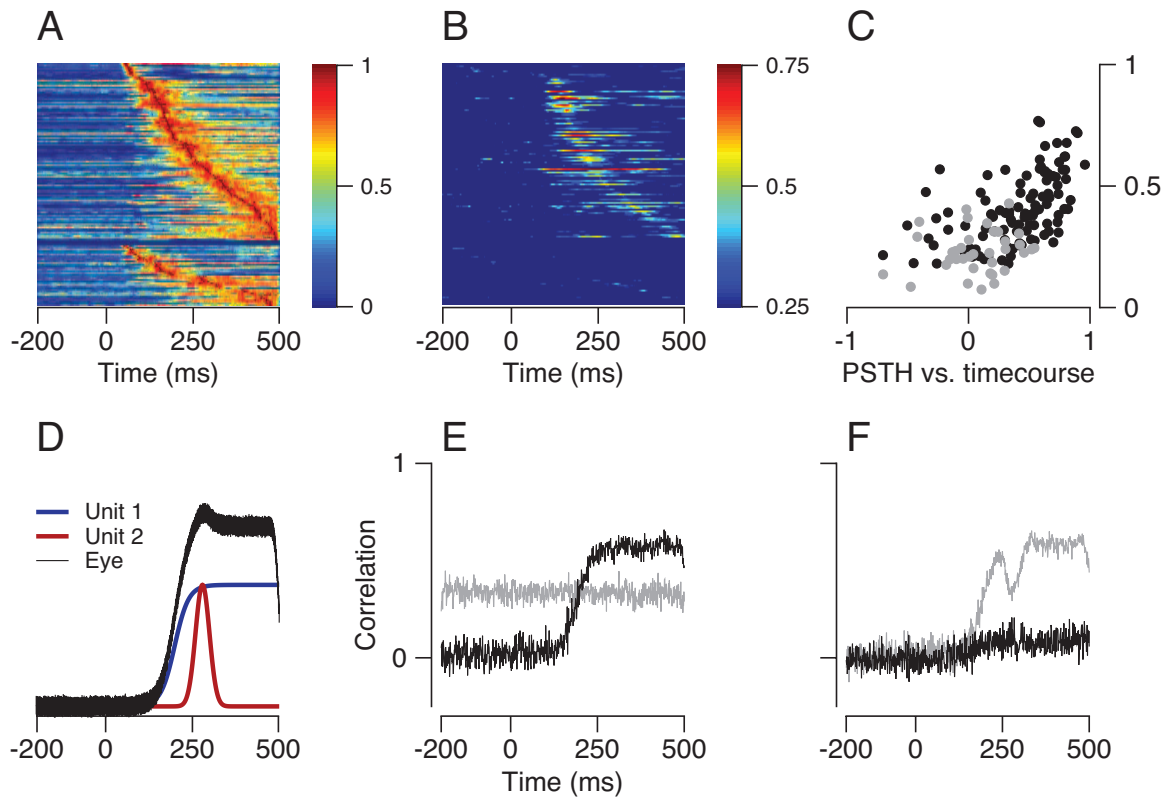


Figure 2.4: Temporal features of population activity, and a model of the consequences of downstream noise. A, The PSTH from each cell was normalized to its maximum activity, and plotted as a color-coded strip. Cells have been ordered according to the latency to peak, and divided into significant (top group) and non-significant (bottom group). B, Each horizontal line of the graph plots the correlation between the actual and predicted residual of eye velocity as a function of time for a single neuron, where color indicate the value of r . Data are grouped as in A. C, The correlation between the PSTH and the time course of neuron-behavior correlations is plotted against the peak value of the correlation. Significant cells are plotted in black, non-significant in grey. D, The filter from two model units is plotted as a blue or red line, and the summed activity from the two (the model eye movement) is plotted in black. E, The time course of the correlation between the signal from the model unit and the summed activity of the two units assuming additive neural noise on the model unit (grey line) and multiplicative noise on the model unit (black line). F, The time course of correlation between the signal from the model unit and the summed activity in the face of uncorrelated multiplicative noise on all units (grey line) and multiplicative noise on both units and their summed activity (black line).

tion between our spikes and the behavior is likely to be quite noisy. In the face of such noise, how could the neuron-behavior correlation come to resemble the PSTH? To address this question, we considered a simple model population. Here, our network is comprised of two independent sub-populations summed together in the face of different types of noise to produce the eye movement. These sub-populations can be thought of simply as two independent neurons, or as the correlated and uncorrelated parts of the network. Figure 2.4D illustrates the time course of activity of the two model sources (blue and red lines), as well as the time course of the command (black) generated by averaging the two traces.

If variations in a neuron's output are additive — that is, if their standard deviation does not depend on the cell's mean firing rate — then the correlation between a given neuron's firing rate residuals and the motor command will be constant as a function of time (Figure 2.4E, gray line). However, if the variations are multiplicative, and scale with the mean firing rate, the correlation as a function of time will closely resemble the sub-population's PSTH (black line), as the fraction of the overall variance contributed by that neuron will rise with its firing rate. Multiplicative variations in firing rate arise naturally from neurons that fire with near-Poisson statistics, such as found in our data set. Therefore, the restricted time course of correlations between neural and behavioral variations may arise in part from the firing patterns of the neurons themselves, in light of the multiplicative nature of their variability.

In contrast, large multiplicative variations in other neurons comprising the active population can cause the neuron-behavior correlation for the neuron of interest to look different from its PSTH. As the variance contributed by these other noise sources rises, the fraction contributed to the eye movement command by the neuron in question will drop. Uncorrelated multiplicative noise in other neurons will tend to remove the inverse of those other neurons' PSTHs from the cell's neuron-behavior correlation, as illustrated by the gray line in Figure 2.4F (see the inverted notch at $t = 280$, the peak of the red cell's PSTH). Multiplicative noise applied directly to the eye movement command will decrease the

magnitude of the correlation, but can preserve the time course of the correlation; such an example is plotted in black.

The close match between the time course of the PSTH and the neuron-behavior correlation in our data implies that multiplicative noise from other neurons or from downstream sources ought to be small — a particularly profound finding, given the synaptic distance of the FEF_{SEM} from the muscles. The fact that the peak of the neuron-behavior correlation predicts the similarity between the time courses can be accounted for by a small population of reliable encoders of movement fluctuations. Our data so far suggests that a small population of neurons, heterogeneous with respect to its strength, features, and time-course of activity could be sufficient to encode the observed movement variability.

2.3.5 An example of the pairwise relationship between filtered spike trains

Thus far, we have dealt exclusively with the relationship between a single neuron and behavior. To further our understanding of the set of neurons responsible for movement variability, we need describe the relationship between the members. One direction of motion from two simultaneously recorded neurons is shown in Figure 2.5. The rasters are plotted in Figure 2.5A, with color-coded spikes; the temporal heterogeneity we have described earlier is clear, with limited overlap in the spikes. Figure 2.5B plots the two filters as a function of lag, as well as a filter derived from shuffled trials, in grey. The red filter is considerably smaller than the blue filter; despite the size difference, the red neuron was over twice as active in this direction: 37.86 (SD: 19.23) spikes/trial, compared to 16.25 (SD: 5.15) spikes/trial from the blue neuron. Figure 2.5C shows the correlation between the predicted and actual behavior as a function of time, as well as the correlation with the filter from randomly shuffled trials (plotted as a gray line).

We've previously established that the neurons in our population are likely contributing to movement variability in a sparse fashion. Therefore, we restrict our analysis to times when both neurons are likely to be part of the active population, as determined by overlap-

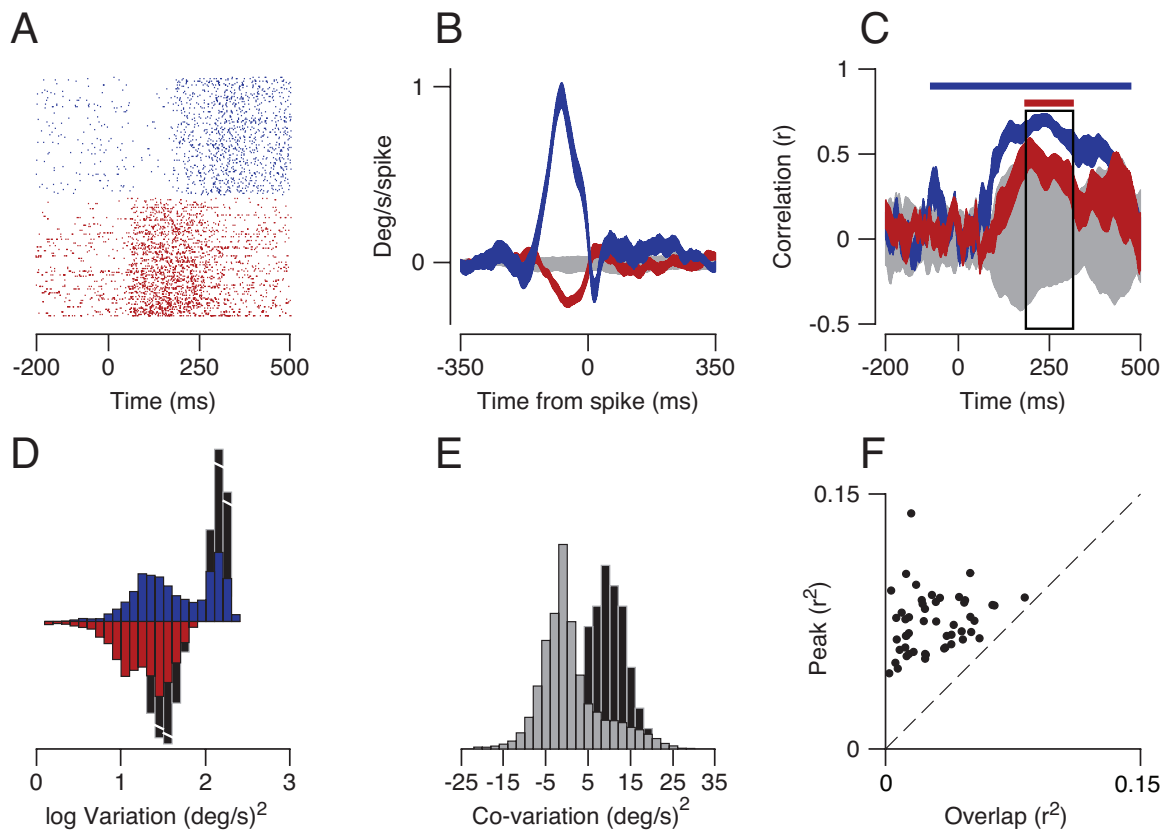


Figure 2.5: The variation, co-variation, and neuron-behavior correlation for an example pair of simultaneously recorded neurons. A, Raster showing the responses of two neurons (blue and red) during the same set of eye movements in one direction. B, The filters associated with each cell, color coded to match; grey lines are the filter derived from shuffled trials. Filter width is the mean \pm 1 SD. C, The time course of neuron-behavior correlation for each unit and randomly shuffled data, color-coded to match. Blue and red horizontal lines show the times when the correlation was 1.5 SD away from the randomly shuffled data. Rectangle shows the time window for analysis, where both cells were significant. D, The variation for all time points are plotted as color-coded bars on a log scale. The red cell has been inverted. Black bars show the corresponding values during the time window of interest, for each cell. E, The co-variation between filtered spike trains for all time points is plotted in grey, and for the time window of interest in black. F, The color-coded distribution of neuron-behavior correlation across the trial, with the time window of interest in black.

ping points in the time-course of the neuron-behavior correlation. Figure 2.5C plots (thick lines) the potential time window for each neurons, with the overlapping region of interest as a rectangle. Within this time region, we examined the variation, and co-variation of the filtered spike trains (the set of predicted eye velocity residuals) from our example pair. The variation is the absolute strength of the fluctuations arising from a set of filtered spike trains. Similarly, the co-variation is the strength of the relationship between the predicted eye velocity residuals from one neuron and its pair. Both quantities are time-vectors, with units of $(\text{deg/s})^2$.

Figure 2.5D plots a histogram representing the probability distribution of the variation on a logarithmic scale ($n = 35,050$). Each neuron is plotted with its respective color, and the sign of the red neuron has been flipped for presentation purposes. The blue neuron has a substantial fraction of values above 10^2 , whereas the red neuron peaks around $10^{1.5}$, confirming that even with fewer spikes, the blue neuron can better describe large fluctuations in behavioral variability, as predicted by size of their respective filters (c.f. Figure 2.5B). The two histograms plotted in black reflect the distribution of values found during the time period of interest, where both neurons encoded behavioral variability well ($n = 6850$). In agreement with our finding that the time course of neuron-behavior correlation tracks the time course of neural activity, the period where the filtered spike trains have their highest values (the tail of the blue and red distribution) corresponds to the period where both neurons have high neuron-behavior correlations.

Figure 2.5E plots the distribution of co-variation between the two units in time, as a grey histogram. Given the opposite signs of the filters (c.f. Figure 2.5B), the large positive tail implies that the red neuron was more active on trials when the blue neuron was less active. The black histogram shows the co-variation present during the time period of interest; it sits well into the tail of the grey histogram, suggesting that there may be a relationship between when the neurons are both actively encoding behavioral variability and when they are most related. Figure 2.5F examines this possibility by plotting the average

correlation (r^2 , so as to value positive and negative correlations equally) during the time of overlap (137 ms) against the average value of the best correlated non-contiguous 137 ms, for each run of the model ($n = 50$). Each point lies above the unity line, demonstrating that our example pair of filtered spike trains are maximally correlation elsewhere than the time of overlap. On average, 24% (SD: 22%) of the highest correlation values were found within the time period when both neurons were well-correlated with the behavior. We conclude that the high values of co-variation found in the time period under examination likely arise from above-average individual activity, and do not reflect a stronger relationship between the two neurons.

We have quantified the relationship between activity from a pair of neurons and ongoing behavior by measuring three quantities: the variation, co-variation, and the neuron-behavior correlation. Further, we have demonstrated the impact of evaluating the data when both neurons can encode the trial-to-trial behavioral dynamics. For this particular pair, the variability is log-normally distributed, suggesting that the encoding capacity of neural fluctuations can be orders of magnitude greater at some times in the trial than in others. Lastly, while the co-variation during the time window of interest is high, it does not reflect the times when this pair of neurons is best correlated.

2.3.6 Quantification of the relationship between members of our population

We turn now to the distributions that characterize our data. 104 pairs of simultaneously recorded units yielded 93 directions with overlapping significant activity, 22 of which had at least 20ms worth of overlapping time; on average 85.18 ms (SD: 56.99). In order to compare equivalently across neurons and between directions, we take the average value across the time window of interest for each direction, for each run ($n = 2200$). Figure 2.6B plots the distribution of neuron-behavior correlations for each neuron within a pair, on each run, for each direction as a set of black bars. Data in Figure 2.6B confirm the veracity of the population maxima plotted in Figure 2.2E, with an average correlation coefficient

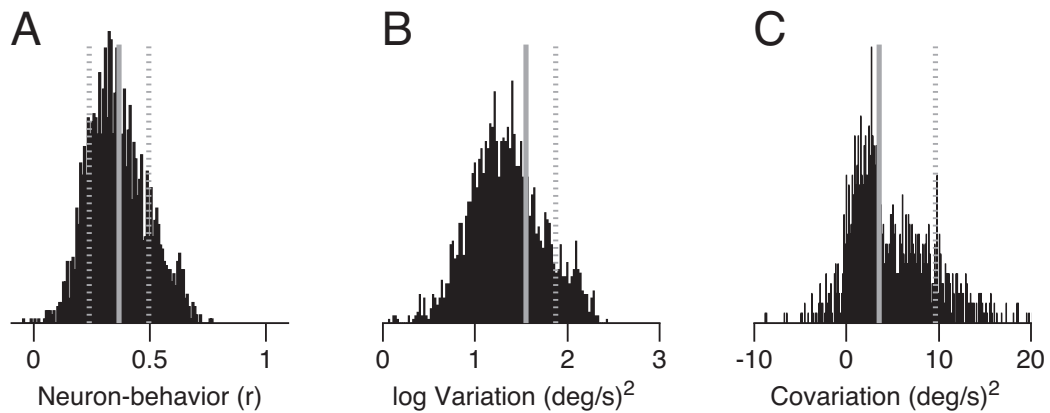


Figure 2.6: A schematic network model, the relevant distributions from all pairs of cells, and a model implementation. A, Each filtered spike trains, schematized by a circle and named with an "X" is summed together to produce an eye movement, schematized by the letter "S." The relevant quantities are the variation ($X_i \cdot X_i$), the co-variation ($X_i \cdot X_j$), and the neuron-behavior correlation, ($X_i \cdot S$). B, The distribution of neuron-behavior correlations plotted as black bars, with the mean and standard deviations in grey. C, The distribution of variation, plotted as black bars on a log scale, with mean and standard deviation in grey. D, The distribution of co-variation, plotted in black, with median and inter-quartile interval plotted in grey. E, The expected value for the neuron-behavior correlations made by the model is plotted against the number of neurons in the active population as a gray line. Error bars represent runs of the model using the standard deviation and inter-quartile interval. Black lines plot the observed mean and standard deviation of the distribution of neuron-behavior correlations.

of 0.37 (SD: 0.13); mean \pm 1SD are plotted in grey. Figure 2.6C plots the distribution of variance on a log scale. The data are well approximated by a log-normal distribution, allowing us to estimate the expected value of, 35.56 (deg/s)^2 (SD: 38.74); expected value + 1 SD are plotted in grey — note that for a log-normal distribution, the expected value is slightly to the right of the mean in log-space. Figure 2.6D plots the distribution of co-variation between pairs of neurons ($n = 1100$). The large concentration of values about zero, combined with the skewed tail and negative values made the distribution difficult to fit, and so we used the median as our estimate of expected value, 3.55 (deg/s)^2 (inter-quartile range: 6.06).

If behavioral variability is generated by correlated activity between large numbers of neurons, then we would expect that when a pair of neurons each individually encodes behavior well, they ought be most correlated with each other. We tested this possibility by comparing the average correlation in time between the filtered spike trains during the time interval of interest to the maximal correlations for the same number of time points. The actual correlation ($n = 1100$, average $r^2 = 0.06$ SD: 0.06) was significantly lower than the maximal correlation values (average $r^2 = 0.07$ SD: 0.06, paired t-test, $p < 10^{-15}$). As with our example pair, there was some overlap between the maximal moments of correlation and the time window of interest, on average 0.45% (SD: 0.36%). Despite the overlap, the difference between the actual and maximum values of correlation leads us to conclude that pairs of neurons were not maximally correlated when they were most likely to jointly contribute to the behavior.

Our population measure of paired data confirms our estimate of the neuron-behavior correlation, as well as the logarithmically distributed nature of the predicted eye movement residuals. We add a measure of the co-variation between individual neurons, and confirm that our pairs are not best correlated with each other when they are both likely to be encoding behavioral variability.

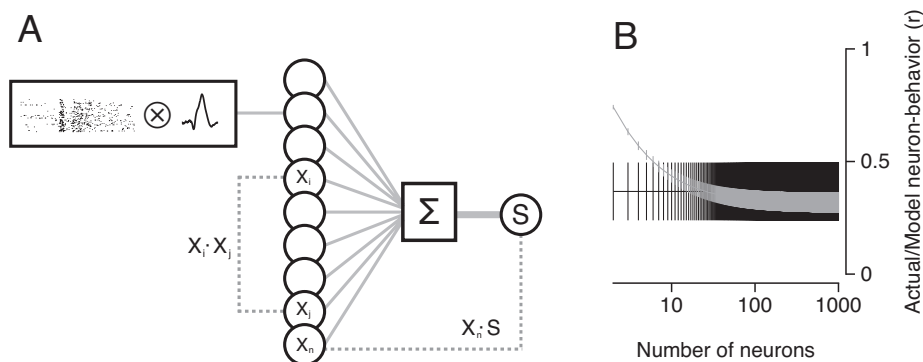


Figure 2.7: A schematic network model, the relevant distributions from all pairs of cells, and a model implementation. A, Each filtered spike trains, schematized by a circle and named with an "X" is summed together to produce an eye movement, schematized by the letter "S." The relevant quantities are the variation ($X_i \cdot X_i$), the co-variation ($X_i \cdot X_j$), and the neuron-behavior correlation, ($X_i \cdot S$). B, The expected value for the neuron-behavior correlations made by the model is plotted against the number of neurons in the active population as a gray line. Error bars represent runs of the model using the standard deviation and inter-quartile interval. Black lines plot the observed mean and standard deviation of the distribution of neuron-behavior correlations.

2.3.7 Our model of the active population is compatible with a sparse set of neurons

With population-level estimates in hand, we can now describe the properties of a network that could give rise to behavioral fluctuations. Our goal is to determine whether a small, sparse population of neurons with properties similar to ours could give rise to the level of correlation between neural and behavioral fluctuations we've observed. Figure 2.7A schematizes our model. On the far left, the rectangle encompasses the drive to behavioral variability: the filtered spike train, with the output drawn as a circle and labeled $X_1 \dots X_n$. The contribution from each neuron is given by the variation ($X_i \cdot X_i$), the co-variation describes the strength of the relationship between elements ($X_i \cdot X_j$), and we postulate that the elements in the set are summed, schematized by the square box, to produce the eye movement, labeled with an S. The neuron-behavior correlation for a given neuron is then $X_i \cdot S$.

At any point in the trial, we estimate a relationship between the size of the population and the expected value of the following three quantities: the variation ($X_i \cdot X_i$), co-variation ($X_i \cdot X_j$), and the neuron-behavior correlation ($X_i \cdot S$). To explain the population activity across the duration of the trial, we make one key assumption: that the expected values of the relevant distributions are stable across time. While the values from different elements may change we assume that the expected values do not, allowing us to model the network once for all time points. Notably, many aspects of our model are equivalent to models used by others to describe the relationship between a population of sensory neurons and a set of perceptual decisions [Zohary et al., 1994, Shadlen et al., 1996].

Figure 2.7B utilizes our model to plot an estimate of the expected value for the neuron-behavior correlation, as a function of the number of neurons in the pool, (grey line). Error bars are given from the SD and inter-quartile range of the relevant distributions. We plot the mean and standard deviation of the observed distribution of neuron-behavior correlations (black error bars). While the asymptote of the model falls below the expected value of the observed neuron-behavior correlations, suggesting the population size falls at their intersection, the variability in our estimates is too high to draw a firm conclusion about the absolute size of the population. Crucially, a network of between 10-30 neurons with properties similar to ours could give rise to the distribution of neuron-behavior correlations we've observed.

Our model is thus compatible with a population where a small handful of neurons (those that are most active) contribute the most to the behavioral drive, and the rest of the population, be it large or small, contributes considerably less. Others have noted that in the presence of correlations, there is little advantage gained in a summative model by increasing the pool size to arbitrarily large numbers; the resulting fidelity of sensory representation remains the same [Zohary et al., 1994]. Similarly, elements of our network remain correlated with the behavior thanks to the correlations, which ensure that the drive from the strongest neurons will be preserved, no matter how many neurons comprise the pop-




Figure 2.8: The filters decorrelate spike trains. A, Black points plot the raw spike-count correlation against the correlation between filtered spike trains. B, The correlation between filtered spike trains is plotted against the distance between the electrodes. Gray points plot the raw data, and black markers the mean and standard deviation for each spacing.

ulation. Thus we conclude that we need not reject the idea our idea of a sparse source for movement variation in the face of co-variation between neurons.

2.3.8 Pairs of filters decorrelate spike trains

Filter shapes can play an important role in determining what happens downstream of the spikes. Correlated spike trains passed through correlated filters retain their correlation, while spikes passing through uncorrelated filters lose theirs. Thus, common input, measured at the level of the spike trains, need not be propagated throughout the network. Our filter shapes form an orthogonal basis set which could remove correlations from the spike trains.

We measured the correlation between filters for the subset of our pairs of neurons that had significant spike-count correlations, evaluated across trials ($n=42/93$, data not shown). Our filters were normally distributed (Kolmogorov-Smirnoff test, $p < 10^{-5}$), mean -0.02 (SD:0.36), implying that pairs of filters were predominantly orthogonal. Figure 2.8A confirms this finding, plotting the raw spike-count correlation against the average correlation between filtered spike trains, for each relevant direction. Only 4/42 points fell above the unity line, demonstrating that the filters largely decrease the level of correlation in the population. The relationship between spike-count correlation and filter similarity just missed

being significant (Spearman's $\rho = 0.30, p = 0.053$), and no relationship remained between the spike-count correlation and the correlation between filtered spike trains (Spearman's $\rho = 0.15, p = 0.33$). Lastly, there was an anatomical basis for the preservation of correlations. Figure 2.8B plots the electrode spacing against the correlation between the filtered spike trains; grey points show the raw data, while black circles show the mean and standard deviation. The relationship was significant, demonstrating that after filtering, spike trains from nearby neurons were likely to be better correlated ($r = -0.55, p < 10^{-3}$).

We conclude that our filters can and do form an effective basis set for removing correlations from common input, as correlated spiking activity is transformed along orthogonal axes, except when the neurons are anatomically close.

2.3.9 Control analyses

In addition to those already described, we performed a set of control analyses. Supplementary Figure 3.3 demonstrates the linearity of the example filter within the space of residuals, the robustness of our fitting procedure, and shows examples of performance on single trials and the distribution across trials. Supplementary Figure 3.4 shows that the filters generalize to both the mean eye movement and off-preferred axes nearly optimally. Supplementary Figure 3.5 shows that that the strength of the predictions was unrelated to any of the following: the neurons' preferred direction; the tuning bandwidth; the average firing rate; the number of trials we were able to record for each direction of target motion; the variance in eye velocity during pursuit; and the variance in the spike train during pursuit. Our controls serve to assure us that our methodology is valid. We therefore conclude that the filters we have described do indeed reflect veridical estimates of what the neurons encode.

2.4 Discussion

2.4.1 Linear systems analysis of continuous oculomotor behavior, across the brain

While many have proposed that cortical neurons, both sensory and motor, are inherently unreliable, our report describes a population of cortical motor neurons whose variation encodes the trial-to-trial deviations in behavior. Neurons in the FEF_{SEM} are thus in agreement with studies suggesting that firing patterns of individual neurons can be quite precise, [Mainen and Sejnowski, 1995, Uzzell and Chichilnisky, 2004] that small numbers of spikes can encode significant information [Luna et al., 2005, Osborne et al., 2004], and that variability in continuous motor behavior is unlikely to reflect true "noise" [Maye et al., 2007]. In sensory systems, it has long been appreciated that quantifying the limits of transmission capability will provide insight into the flow of information between neurons. By demonstrating that single spikes can encode information about behavior, we can extend the analyses established in the sensory domain into the realm of movement.

Linear systems analysis of motor systems offers a number of special challenges. The method we used to analyze the co-variation of neural firing and behavior is adapted from techniques traditionally used to describe the relationship between neural firing and a sensory stimulus [Chichilnisky, 2001, Nagel and Doupe, 2006, Wu et al., 2006, Dayan and Abbott, 2005, Rieke et al., 1999]. Because of the slow nature of smooth eye movements, our filters are probably wider than the actual influence of each cell on its immediate downstream targets, even though we corrected for the first-order autocorrelations. The filters accurately approximate the relationship between spikes and the real smooth eye movement, and should be thought of not as models of FEF_{SEM} neurons *per se*, but rather of the entire downstream cascade that follows an FEF_{SEM} spike.

When choosing to analyze the variation in a set of closely related motor behaviors — stereotyped responses to an identical stimulus — we hoped to find a range of motor out-

puts over which the neural response would be largely linear. There may yet be a nonlinear model relating neural spiking to eye movements that would reveal even greater correlations. To our knowledge, all relevant non-linear models perform better than their linear counterparts, confirming our belief that our method represents a lower bound on the possible correlations between neural and behavioral residuals. Perhaps the most promising way to extend our model is to adopt an approach that has proven successful in the sensory domain: to model the transformations that occur throughout the sensory-motor circuit as a cascade of linear/non-linear systems [Rust et al., 2006]. Additionally, it may even be possible to exploit new methods for estimation of non-linear filters and come closer to deriving a near-optimal filter for a given data set [Pillow and Simoncelli, 2006, Sharpee et al., 2004]. By similarly measuring the transformation between neurons and behavior at the motor periphery, and moving upstream, we may be able to gather the data necessary to extend our model of FEF_{SEM} neurons to explain the transformations between areas, and not simply between one area and ongoing behavior. We can extend this approach further to quantify the relationship between each transformative stage along the sensorimotor continuum, to ultimately characterize how populations of cells control behavior.

2.4.2 A sparse network model in the face of co-variation

Our analysis of the correlation between neural and behavioral variability demonstrates that "extra" spikes in single neurons in the FEF_{SEM} are strongly correlated with subtle trial-to-trial variations in behavior. The strongest form of our hypothesis suggests that probabilistically speaking, in a brief time window, a single additional spike in a given neuron will lead to a deviation of known shape and size in the eye movement. If true, what would this tell us about the network of cells that generate movement variability?

One architecture such a network might have is a set of independent neurons, each making their own unique contribution to movement; under such circumstances, a sparse source for movement variability is quite reasonable. However, the presence of correlated

activity between pairs of cells allows us to rule out such a network. Alternatively, the network of cells might be both large and correlated, with the activity from many units ultimately averaged together [Zohary et al., 1994, Shadlen et al., 1996]. In such a network, movement variability would arise from common aspects of variation, which is all that could be preserved in the face of downstream aggregation. Could such a model of common drive explain our observations?

The orthogonality of our filters, and the ensuing decorrelation of spike trains suggests that common fluctuations in the set of neurons we observed are unlikely to be the sole source of movement variability. To preserve the common drive hypothesis, we must subdivide our set of neurons by their filter shape, thus allowing neurons to have differential access to behavior. Our data suggests that neurons encode behavioral fluctuations at different moments, and so we must further aggregate our population into groups with similar filters, temporal profiles, and preferred direction. In light of the many aspects of oculomotor behavior our task has not addressed, such as speed, vergence, and vestibular sensitivity, the behavioral space rapidly becomes quite large.

Ultimately, though, our data are not compatible with the idea that common drive from large populations of neurons in the FEF_{SEM} is the source of behavioral fluctuations. Both the magnitude of the filters — our best estimate of the consequences of a single spike — and our measurement of variation in behavioral drive, are log-normally distributed. The implication of such a distribution is that some neural fluctuations are orders of magnitude better at encoding behavioral variability than others. In the framework of the "common drive" hypothesis, cells that best encode behavioral fluctuations must be cells that encode common elements of population fluctuations. Mechanistically, the most parsimonious explanation for the presence of such cells is that they are the source of correlated neural variability. This is not to imply that a handful of cells have a preferential ability to drive behavior. Instead our data is compatible with the idea that variability in certain cells, at certain moments, on certain trials, are preferentially amplified to produce behavioral fluctu-

tuations.

For technical reasons, we restricted our analysis of the encoding capacity of single neurons to the domain of movement variability, as opposed to the totality of movement. Analytically, filters derived from the residuals are the same as derived from the complete movement; our control analyses of filter generality support this notion. We therefore feel that our conclusions about a sparse origin of movement variability could well extend to the population of neurons that give rise to the totality of behavior.

2.4.3 The role of spikes in the FEF_{SEM}

Are single spikes capable of eliciting behavior? Invertebrate neurobiology has identified "command neurons" that actuate ethologically relevant behavioral cascades; many quite complicated [Pearson, 1993]. Furthermore, recent evidence shows that rodent whisking can be modulated by single spikes, suggesting that even in mammals, single neurons can affect motor behavior [Brecht et al., 2004]. Does the strong relationship we observe suggest a similarly strong causal role for spikes in the FEF_{SEM} ? Both microstimulation studies [Bruce et al., 1985, Gottlieb et al., 1993] and lesion studies [Keating, 1991] have described the causal role the FEF_{SEM} plays in determining the parameters of eye movements: while not unique in this respect, activity in the FEF_{SEM} is likely necessary for normal pursuit dynamics.

The strong correlations observed between neurons and behavior might reflect a highly precise command for movement, or an "efference copy" of the ongoing movement. Such a representation could arise from the corollary discharge inputs that may well contribute to the responses of neurons in the FEF_{SEM} [Wurtz and Sommer, 2004]. The latencies of our filters do not resolve this issue, partly because the lag from the spike to the time of maximum influence is broadly distributed around zero, and partly because the autocorrelations in the pursuit response limit our ability to resolve the precise latency of the neuron-behavior correlations. Therefore, we cannot exclude the possibility that the neurons we've recorded

from play no causal role in determining behavior, even though they are high-fidelity encoders of the movement, a finding that would be fascinating in its own right. However, we think that the other evidence of a causal role for neurons in the FEF_{SEM} renders this conclusion unlikely.

2.4.4 A sparse temporal code for movement variability in the cortex

Our working hypothesis to explain the temporal diversity of neuron-behavior correlations, and its relationship to the temporally restricted neural activity, is that for each period of a trial, the natural variability in a small sub-population of FEF_{SEM} neurons gives rise to the fluctuations observed in the eye movement. A sparse source for movement variation has a number of advantages. First, it would allow the cortex to exercise modulation and voluntary control over different aspects of motor behavior on a fine temporal scale, allowing different neurons that mediate pursuit to have special influences at different times; the orthogonal nature of the filters supports this notion. Second, it would enable excellent temporal control over long-term plasticity and learning (Fiete et al., 2004). Available data imply that the site of pursuit learning is downstream from the signals provided by the FEF_{SEM}: recordings from the FEF_{SEM} failed to reveal any changes that accompanied pursuit learning, while learning did cause changes in the smooth eye movements evoked by stimulation in the FEF_{SEM} [Chou and Lisberger, 2004]. Pursuit learning is exquisitely sensitive to time, and both the direction of target motion and the time when it changed can be learned well [Carey et al., 2005, Medina et al., 2005]. A temporally sparse representation of pursuit in the FEF_{SEM} could thus be an important factor underlying temporally specific learning at a downstream site such as the cerebellum. We propose that the neurons that can encode movement with high-fidelity are likely to be preferentially involved in learning. Further, we note that changing a small set of neurons is likely to be more efficient than synchronously modulating thousands, and may underlie our remarkable ability to rapidly learn new motor behaviors while still retaining old ones.

Recent work in motor control in songbirds and representation of visual stimuli provides evidence that ethologically relevant stimuli and behaviors can be quite sparsely represented, with remarkable sensitivity and precision; here we provide complementary evidence for sparse coding in a similarly relevant behavioral space [Hahnloser et al., 2002, Theunissen, 2003, Vinje and Gallant, 2000]. Since smooth pursuit is considerably slower than most behaviors, sparseness in oculomotor neurons ought be comparably restricted. We feel that the idea of sparse network organization for the generation of behavior can accommodate many behaviors, no matter their specific coding demands.

2.4.5 The origin of behavioral fluctuations

In light of the temporal similarity between neural activity and neuron-behavior correlations, we have suggested that the noise downstream of the FEF_{SEM} must be quite low. The finding is most remarkable in light of the number of synapses between the FEF_{SEM} and the muscles, and demands high-fidelity transmission of information downstream. Naturally, a reliable motor system is compatible with noise originating in the FEF_{SEM} itself, or sensory areas upstream, in agreement with behavioral experiments from our group [Osborne et al., 2005, 2007].

Evidence for reliable transmission of variability is found in our analysis of the output of the cerebellum (Medina and Lisberger, under review). In contrast to the temporal sparseness of the FEF_{SEM} , responses of Purkinje cells in the flocculus are temporally stereotyped across the population. In excellent agreement with the attendant neuron-behavior correlations, neural activity peaks during the initiation of pursuit and is sustained at a lower level during steady-state tracking, and thus furthers our proposition of limited downstream noise. Further, Purkinje cells are largely tuned to cardinal directions — consonant with the pairs of muscles that drive the eyes — while neurons in the FEF_{SEM} are tuned across all possible eye movement directions. The stereotyped responses of the Purkinje cells may serve to transforming the spatially-oriented movement commands into coordinates appro-

appropriate for the muscles in a high-fidelity manner. The cortical and cerebellar data are thus in agreement with the notion of a sensory origin for motor noise.

2.4.6 What is a single spike worth?

In this paper, we show a set of data, analyses and models compatible with the notion that variability in single cortical neurons can give rise to movement variability. While we cannot yet watch the entire population of neurons involved in generating smooth pursuit eye movements, it is this very limitation that makes understanding the nature of the active population intriguing and controversial. We hope that this report serves to inform this most interesting debate, and appreciate that it is far from the final word.

2.5 Methods

Two male monkeys (*Macaca mulatta*) were used in experiments approved by the Institutional Animal Care and Use Committee of the University of California, San Francisco. All experimental procedures were in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals. Eye position was monitored using the scleral search coil technique as described previously while the head was held stationary using custom hardware [Ramachandran and Lisberger, 2005]. The eye coil and hardware were previously implanted during sterile surgery with the monkey under Isoflurane anesthesia. One monkey was an experienced subject, having previously been well-trained to smoothly track visual targets for a juice reward; he needed no additional behavioral training for this study. The other monkey was naive to experimental methods, and was trained over a period of one month until his performance was indistinguishable from other, expert, animals.

We used traditional step-ramp target motion to elicit excellent smooth pursuit responses from our subjects. A small spot of light was presented at the center of an analog oscilloscope screen (14 inches diagonal), and after a variable fixation period the light stepped and

ramped in opposite directions such that the target would cross the original position of fixation at approximately the time of maximal smooth eye velocity. Targets always moved at 20 deg/s in one of eight evenly spaced directions. Monkeys had to maintain their gaze within a small window (0.75 degrees during fixation, and 2-3 degrees during target movement) to receive a small fluid reward. The duration of monkey Cb's pursuit started between 500-700 ms, and was increased for the last $\sim 2/3$ of experiments, first to 750-950 ms and then to 900-1100 ms, to gather data over longer periods of sustained behavior. Monkey Gu always tracked the target for between 900-1100 ms.

Up to five quartz-shielded tungsten electrodes were lowered simultaneously into the caudal region of the frontal eye fields using a microdrive (Thomas Mini-Matrix 05, Thomas Recording, Giessen, Germany). When feasible, we sought to record from multiple single units simultaneously, both on the same and on different electrodes. Electrical signals were amplified and digitized for on-line spike sorting using template matching of waveforms that crossed a user-defined threshold (Plexon MAP, Plexon Inc., Dallas TX). Waveforms were considered to reflect the activity of a single source if they were distinct from the noise and had no refractory period violations (more than 2 threshold crossings within 1ms). Additionally, during offline sorting, when plotted as projections onto the first and second principal components, units had to form clusters distinct from both noise and other units over the duration of the experiment (usually > 45 minutes of continuous data acquisition). After sorting, waveforms were converted to timestamps with 1 ms precision for analysis.

We studied the smooth component of the monkey's pursuit of a moving target. The movement period was defined as 200 ms of fixation before the onset of target motion, and 500 ms after the target began to move. Velocity traces were generated by digitally differentiating the position traces with a balanced difference algorithm; traces were digitally smoothed with a zero-phase 25 Hz 2-pole Butterworth filter. Saccades show up as large deflections in the velocity trace and were removed so as not to dominate the trial-to-trial variations: saccades were marked initially with a threshold-crossing algorithm, described

previously, and each trial was examined by eye using custom software [Schoppik and Lisberger, 2006]. The velocity trace then was interpolated between saccade beginning and end. We restricted our analysis of the relationship between behavioral and neural variability to the eye movements along the single best direction of the neuron, determined by the direction that had the greatest number of spikes overall. When the preferred direction was along an oblique axis, we analyzed the radial eye velocity. There was little jitter in the time of movement onset, so we defined the appearance of the moving target as the beginning of the behavior.

Each neuron's raster was examined by eye, to exclude units that were not strongly responsive during the eye movement, or whose responses changed qualitatively over the duration of the experiment. A minimum of 15 repeats of pursuit in a single direction was necessary for inclusion; on average, we had 80 (SD: 53) repetitions of each trial in our data sets. A total of 141 single units met our criteria for inclusion in the analysis (97 and 44 from monkeys Cb and Gu). As we could not resolve qualitative or quantitative differences between the data from the two monkeys, all data were analyzed as a unified population.

Although a smoothed estimate of firing rate was not used for these analyses, we did need to generate one for the schematic shown in Figure 2.2A. To obtain a continuous estimate of firing rate for our diagrams, we smoothed the binary spike train for that experiment using a Gaussian filter with a standard deviation of 20ms. Similarly, the estimate of the PSTH in Figure 2.4A was determined in a 10ms sliding window.

2.5.1 Modeling the relationship between the residuals in spikes and eye movements

Our analysis consisted of deriving a linear filter that related the variation in neural firing to the variation in eye movement responses on a trial-by-trial basis. The residual variation was generated by subtracting the mean eye trajectory across all trials from the eye's trajectory on each trial, and similarly for the spike train. Filters were then generated by

cross-correlating the residuals of firing rate and eye velocity on each trial, and taking the average of across trials. To avoid over-fitting, we used a cross-validation technique that randomly selected a subset of data to generate the filter, and used the remaining trials to test the performance of the filter. The split that minimized both the variance in the estimate of the filter and the variance in the prediction quality was 40% filter generation and 60% tests (Supplementary Figure 8). To generate confidence intervals on the filter, we repeated the reconstruction procedure 150 times with different trials assigned for filter generation or testing. To estimate the null hypothesis, that there was no relationship between trial-by-trial neural and behavioral variation, we generated a separate filter after shuffling the neural responses across trials. This procedure is conceptually similar to the "shuffle" or "shift-predictor" used to estimate the effect of a repeated stimulus on correlations between neurons [Brody, 1999].

Specifically, we began with the following model of the transfer function between eye movement residuals and spikes:

$$e_i(t) - \bar{e}(t) = D(\tau) \otimes (\rho_i(t) - \bar{\rho}(t)) \quad (2.1)$$

where $e_i(t)$ and $\rho_i(t)$ are the eye speed and the spike train on trial i , and $\bar{e}(t)$ and $\bar{\rho}(t)$ are the mean eye movement and mean neural response, respectively. Our model assumes that there exists some linear filter, $D(\tau)$ that allows us to convert the residual spike variability into the residual eye movement. To estimate $D(\tau)$, we used the following equation:

$$Q_{e\rho}(t) = \frac{1}{n} \sum_{i=1}^n [(e_i(t) - \bar{e}(t)) \otimes (\rho_i(t) - \bar{\rho}(t))] \quad (2.2)$$

$Q_{e\rho}$ is the filter we have generated, and is some approximation of the true filter $D(\tau)$. $Q_{e\rho}$ differs from $D(\tau)$ by the autocorrelation of the spike train [Theunissen et al., 2000]. To correct for correlations in the spike train, we divided the filter in the frequency domain by the power spectrum of the residual spike train:

$$\hat{D}(\tau) = \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega \frac{\sum_{i=1}^n Q_{e_i\rho_i}(\omega)}{\sum_{i=1}^n Q_{\rho_i\rho_i}(\omega)} \quad (2.3)$$

where $\hat{D}(\tau)$ is the decorrelated filter, and $Q_{e_i\rho_i}(\omega)$ and $Q_{\rho_i\rho_i}(\omega)$ are the Fourier domain representations of the filter from Equation 2.2 and the power spectrum of the spike train, respectively. Equation 2.3 states that the true eye movement reconstruction filter can be estimated by weighting the power in the raw filter by the power in the spike train. Because the spike train has a relatively narrow autocorrelogram, the majority of its power lies above 50 Hz. Consequently, the denominator of Equation 2.3 acts as a low-pass filter.

By convention, we call the filter we have generated the "reverse" filter, as it allows us to predict eye movements using spike trains. To generate the "forward" filter, which predicts spike trains from eye movements, we must correct for the autocorrelations in the eye movement. We replaced the denominator in Equation 2.3 with the power spectrum of the residual eye movements. Noise in the estimate of the power spectrum tends to overwhelm the signal in the filter, and so we smoothed the spectrum using an exponential filter with a rapid dropoff (~ 2.5 Hz). The cutoff for smoothing was set at the frequency where the power in the spectrum of the shuffled residual eye movements (our estimate of a noise spectrum) overlapped the power in the spectrum of the residual eye movements. The mean dot-product between the forward filter flipped about the zero-lag point and the reverse filter was 0.91 (SD: 0.05), supporting the conclusion that despite a different denominator, the two filters ought be equivalent. We thus focus the description of our results on the "reverse" filter, which needed no additional smoothing. Lastly, we computed the coherence of the filter:

$$C(\omega) = \frac{Q_{ep}(\omega)^2}{Q_{ee}(\omega) * Q_{pp}(\omega)} \quad (2.4)$$

where Q_{xx} is iconic for $\sum_{i=1}^n Q_{x_i x_i}(\omega)$. Coherence describes the percent of variance in the data accounted for as a function of frequency, and thus is similar to the definition of r^2 , which is the square of Pearson's correlation coefficient. Importantly, with two terms

in the denominator, the coherence corrects for autocorrelations in both the eye movement residuals and the spike train residuals.

We defined a filter as "significant" if a two-sided t-test showed that the total power in the filter was significantly different from the total power in the filter derived from randomly shuffled data. The value of α was adjusted by the number of draws, so that our criterion for significance became $p < 0.0003$. In monkeys Cb and Gu, 74/97 and 31/44 neurons yielded filters that were significant, respectively. By convolving the successfully derived filters with the residuals of the spike trains on each test trial, we generated predictions that could be compared directly with the residual eye velocity on each trial. We measured the model performance by calculating Pearson's correlation coefficient for the predicted vs. actual residuals, point-by-point. To estimate the significance of the predicted residuals, we shuffled the eye velocity and firing rate residuals and recalculated the correlation coefficient on shuffled trials. All significant filters made statistically significant predictions.

2.5.2 Data and analysis of paired recordings

Our data set included 104 simultaneously recorded pairs of neurons, (84 and 20 from monkeys Cb and Gu), with 26 pairs recorded on the same electrode. As the neurons rarely had overlapping preferred directions, we had to extend our analysis to the complete set of eight directions of motion. We computed the filter as in Equation 2.3, direction by direction, for each neuron in the pair, for 50 runs in each direction. Covariation and correlation were only evaluated for directions where both filters were significantly different from shuffled trials, ($n = 43/104$ pairs with overlapping significant filters, $n = 93$ directions from the 43 pairs)

Our model for the relationship between brain and behavior, given by Equation 2.1, states that the filter convolved with the spike train residuals is our prediction for the eye movement residuals on that trial. We used this to model the relationship among a pop-

ulation of filtered spike trains, estimating three relevant quantities. The neuron-behavior correlation as a function of time has been defined earlier. The variation and co-variation are used to estimate the parameters of the covariance matrix, and are defined according to the following equation:

$$F_{ij}(t) = [D_a(\tau) \otimes (\rho_a(t) - \bar{\rho}_a(t)) \cdot D_b(\tau) \otimes (\rho_b(t) - \bar{\rho}_b(t))] \quad (2.5)$$

where $D_n(\tau)$ represents the filter from a given neuron, n , and $(\rho_n(t) - \bar{\rho}_n(t))$ represents the matrix of spike train residuals from neuron n . The equation thus compares the relationship between the filtered spike trains. By aligning the matrix appropriately, we could evaluate the relationship across trials, producing a vector of values in time. Alternatively, we could evaluate the dot-product aligned such as to compare the trials across time, producing a vector as long as the number of trials. By definition, if the two matrices of filtered spike trains were identical ($a = b$), Equation 2.5 yields the variation, and if different ($a \neq b$), the co-variation. To evaluate the correlation, we added a denominator to Equation 2.5 consisting of the square root of the product of the variance in each matrix.

To analyze the relationship between spike trains before and after filtration, we needed a way to compute the correlation across trials. We measured the correlation between spike trains across the trial by computing the familiar spike-count correlation (rSC), defined as the correlation between a vector containing the number of spikes on each trial (evaluated during pursuit) for each neuron. Statistical significance was evaluated by comparing the spike-count correlation to the correlation derived from randomly shuffled trials. The rSC was computed only on the 60% of trials evaluated on each run of the model, allowing us to compare directly to the correlation between filtered spike trains, evaluated using equation 2.5, taken across time, rather than trials.

We used standard statistical methodology to describe our data. Correlation coefficients were evaluated using Pearson's method; when the data contained large outliers, we used

Spearman's ρ . To compare normally distributed data, we used t-tests, paired when appropriate; non-normally distributed data, were compared with the Wilcoxon rank-sum test. Finally, circular correlation was evaluated using Mardia's r test. All code was written using MATLAB (Mathworks, Natick MA), and is available, along with all relevant data, at <http://keck.ucsf.edu/~schoppik/>.

Chapter 3

Supplementary Material

3.1 Control Analyses

3.1.1 Abstract

In addition to the analyses described in 2, we provide a number of complementary investigations to examine the validity and generality of our filters. We start by examining the performance of linear filters derived on the position trace, as well as on eye movements aligned to minimize variability. Next, we examine the consequences of changing the size of the bins we use to examine our data. We then present a complete analysis of the predictions made by our example neuron. We describe how our filters perform when asked to evaluate the mean eye movement, as well as the eye movement along directions other than the preferred. We present data that show that the performance of our filter is not related to common parameters of oculomotor performance. We use PCA analysis to confirm our clustering of the filters, and finally, we justify the fractions of data we chose for cross-validation runs of our model.




Figure 3.1: The relationship between spikes and behavior does not depend upon the representation of the eye movement, the trial alignment, or the trial length. Each plot shows the peak correlation between actual and predicted residuals of eye velocity for each cell in the population under different analysis conditions as a function of the peak correlation derived from the eye velocity traces used throughout the paper. A, filters based on the eye position trace, including saccades, B, filters based on eye velocity traces that had been aligned in time so as to minimize the overall variation among the responses, and C, filters based on eye velocity traces that were shortened in duration. X'es and gray circles show data from neurons with statistically-significant and non-significant filters. The unity line is plotted as a dashed black line.

3.1.2 Our findings obtain in the position domain, as well as for aligned data

How important was our choice to represent behavior as the eye's velocity? We opted to use the eyes' velocity, after saccades had been removed because catch-up saccades, being almost an order of magnitude faster than normal pursuit dominate the trial-to-trial fluctuations. Additionally, neurons in the FEF_{SEM} are not thought to be sensitive to saccadic eye movements, or the eyes' position in the orbit, and so we opted to exclude information about the eyes' position [Tanaka and Lisberger, 2002b]. To confirm that our findings were not limited to the particular representation of the eye movements we chose, we re-ran the analysis using the complete eye position trace. Supplementary Figure 3.1A plots the outcome of analyzing position against the analysis of velocity. The strength of the relationship between spikes and behavior for each cell is quantified by the peak of the correlation between the predicted and actual variability, measured in time. Cells that yielded a significant filter are marked with a black "x," while non-significant filters are shown as gray dots. The cells continued to do well at predicting the variability in the position trace, as marked by a mean correlation for the significant filters of 0.38 (SD: 0.17). The predictions made by significant filters generated in the position or velocity domain are not distinct (paired t-test, $p > 0.15$). We propose that cells performing better using the position analysis are sensitive to variation either in the position, or at the high frequencies present during the saccade. Similarly, cells that perform better using the velocity trace are likely sensitive to the slower changes in the eyes' movement present during smooth pursuit. We conclude that our choice of the velocity trace, with saccades removed, represents an appropriate measure of the behavior.

Our method of filter generation relies upon an accurate estimate of the variation. In our analysis, we have aligned the data at the start of target motion, rather than at the moment that the eye begins to move. If the spike trains carry information that is time-locked to the start of the eye's movement, then this will appear as "variation" in our analysis. Such precise timing from trial-to-trial is interesting in its own right [DiCarlo and Maunsell, 2005,

Pouget et al., 2005], but we were concerned that it might lead to an overestimation of our models' performance. Rather than re-aligning the eye movement traces according to some arbitrary measure of the latency, we opted to simply eliminate as much of the trial-to-trial variability as was possible through temporal realignment.

We realigned our trials by shifting each trial so as to minimize the in a least-squared distance between all eye velocity traces in a given experiment. This had the effect of drastically homogenizing the trials in time. As a consequence of realignment, our trials began and ended at different points: we shortened the trials appropriately so that they were all of the same duration. Supplementary Figure 3.1B plots the outcome of our analysis using the aligned traces against the outcome of the analysis on unaligned data. The strength of the relationship between spikes and behavior for each cell is quantified by the peak of the correlation between the predicted and actual variability, measured in time. Cells that yielded a significant filter are marked with a black "x," while non-significant filters are shown as gray dots. Two features bear highlighting: first, despite removing as much of the variation from trial-to-trial that we could by realignment, cells continued to do well: the mean correlation between actual and predicted residuals of eye velocity was 0.33 (SD: 0.14) for significant filters, versus 0.41 (0.15). Secondly, the majority of the points fall below the unity line, suggesting that the filters do better on unaligned data; the aligned data performed significantly worse (paired t-test, $p < 10^{-6}$). We conclude that the variability introduced by our choice of alignment does contribute to our estimate of filter performance; but our main findings persist even when the trials are optimally aligned.

Our realignment was performed on shorter trials; the missing data may have contributed to the slight underperformance of aligned trials. We repeated our analysis on trials aligned to the onset of target motion as in the Results, but that had been truncated to match the length of the re-aligned trials. Supplementary Figure 3.1C plots the outcome of our analysis using the shortened traces against the data described in the paper. Each point plots the peak of the function describing the correlation between the predicted and the

actual eye variation as a function of time. Cells that yielded a significant filter are marked with a black "x," while non-significant filters are shown as gray dots. The high correlation between the two data sets demonstrates that the findings described in Supplementary Figure 3.1B are not due to shorter trials ($r = 0.92$, $p < 10^{-43}$). The mean of the significant filters generated on shorter trials was slightly lower, 0.39 (SD: 0.15) compared to 0.41 (SD: 0.14) (paired t-test, $p < 0.01$). More importantly for the analysis, the data in Supplementary Figure 3.1C show that our finding is robust to losing 10-15% of the trial duration. Beyond this point, though, the data set is likely to be too short to describe the entirety of the filter, given the distribution of filter widths and latencies that we have observed.

3.1.3 Increasing bin size decreases our signal-to-noise ratio

In our analysis, we did not bin our data beyond our initial sampling rate. Many studies of neural activity present data that has been smoothed in some fashion, by binning, windowing, or by some more sophisticated method. The signal to noise ratio of many analyses can improve with increasing bin size, and at first glance, the same should hold true for our analysis: we derive the transfer function from the residual eye movements and spike trains, a linear operation. Binning too is a linear operation, and thus we should be able to filter our data before or after derivation of the transfer function, with an accompanying improvement in the signal to noise ratio. However, division by the power spectrum of the spike train (Equation 2.2) introduces a non-linearity in filter generation. This non-linearity is the source of a counter-intuitive consequence of binning data: the less binning, the smoother our estimate of the transfer function — our behavior is smooth and continuous, and so a smoother filter can better capture the relevant frequencies. Supplementary Figure 3.2 demonstrates the consequences of binning on the smoothness of the resulting filters.

Supplementary Figures 3.2A and 3.2D show the residuals for the spike train binned into 1 ms or 20 ms bins. Visually, they differ because the absolute number of spikes/bin

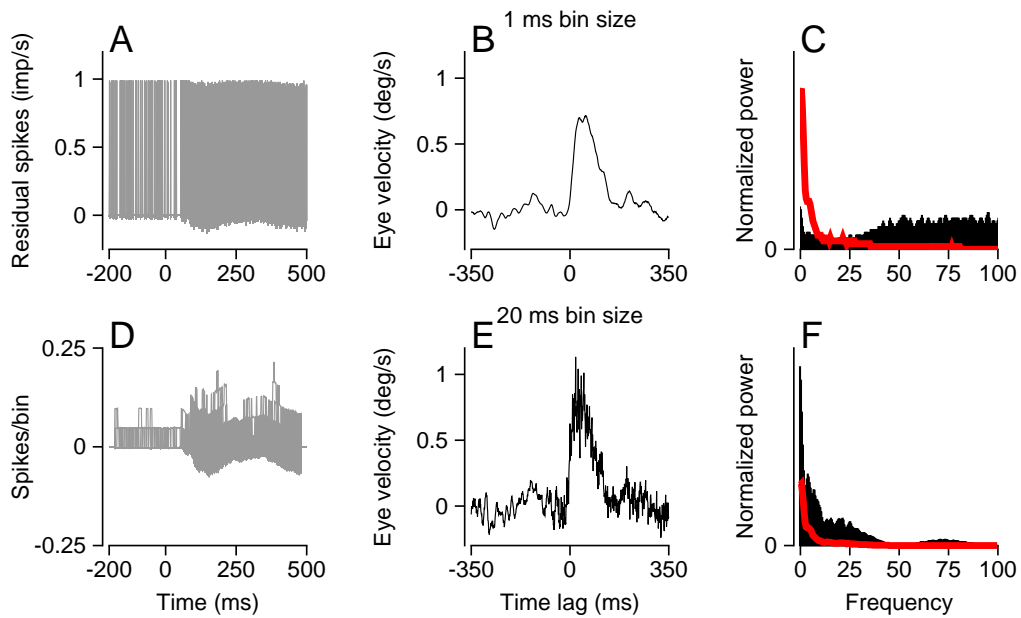


Figure 3.2: Smaller bin sizes lead to lower noise in the estimate of the relationship between spikes and eye movements. A,D, The variation in the spike train, evaluated at 1 ms bin size and 20 ms bin size. B,E, The filter generated from the spike data in A,D. C,F. The normalized power spectrum of the residual spike trains is plotted in black for 1 ms and 20 ms bin sizes. The normalized power spectrum of the filter in B,E, is plotted in red.

increases with the size of the bin. The resulting filters in Supplementary Figures 3.2B and 3.2E have similar shapes, but differ in that the use of a 20-ms bin leads to more high frequency noise in the filter. Here, the filters differ because they have been "decorrelated" by the power spectrum of the spike train, and the power spectrum of the spike train depends on the bin width. Recall from Equation 2.4 that the power spectrum of the spike train is in the denominator while that of the raw, un-normalized filter (Supplementary Figures 3.2C and 3.2F, red curve) is in the numerator. When the bin width is 1 ms, there is little power in the spike train variability at low frequencies but considerable power at high frequencies: the denominator acts as a low-pass filter. When the bin width is 20 ms, in contrast, the spike train has more power at low frequencies and less power at high frequencies: as a result, the signal to noise ratio of the filter is decreased substantially. Practically speaking, then, the finer the resolution used to generate filters, the smoother the filter.

3.1.4 A detailed evaluation of the predictions made by the example neuron

How well do our predictions fit for a single neuron? In 2 we describe the performance of the filter in time, but there are many other ways to describe the results of our analysis. Here we describe the complete predictions, point by point, as well as trial-by-trial. Our techniques for quantifying the strength of the transformation between brain and behavior assume the relationship is linear. By choosing to model the filter between the residual spike train and the residual eye movement, we hoped to constrain our analysis to a regime where the relationship was likely to be linear. To examine the behavior of our filter, we combined the data from all trials and plotted the predicted versus actual eye movement residuals for each time point in each trial (Supplementary Figure 3.3A). The relationship can be well-fit by a line, shown both in the raw samples (gray symbols) and in the means obtained by averaging within 20 bins chosen so that they contained equal numbers of points (black symbols and line). Many points cluster around 0 because they included data from 200 ms of fixation, corresponding to the mottled region in the upper left corner of Figure 2.2A, and

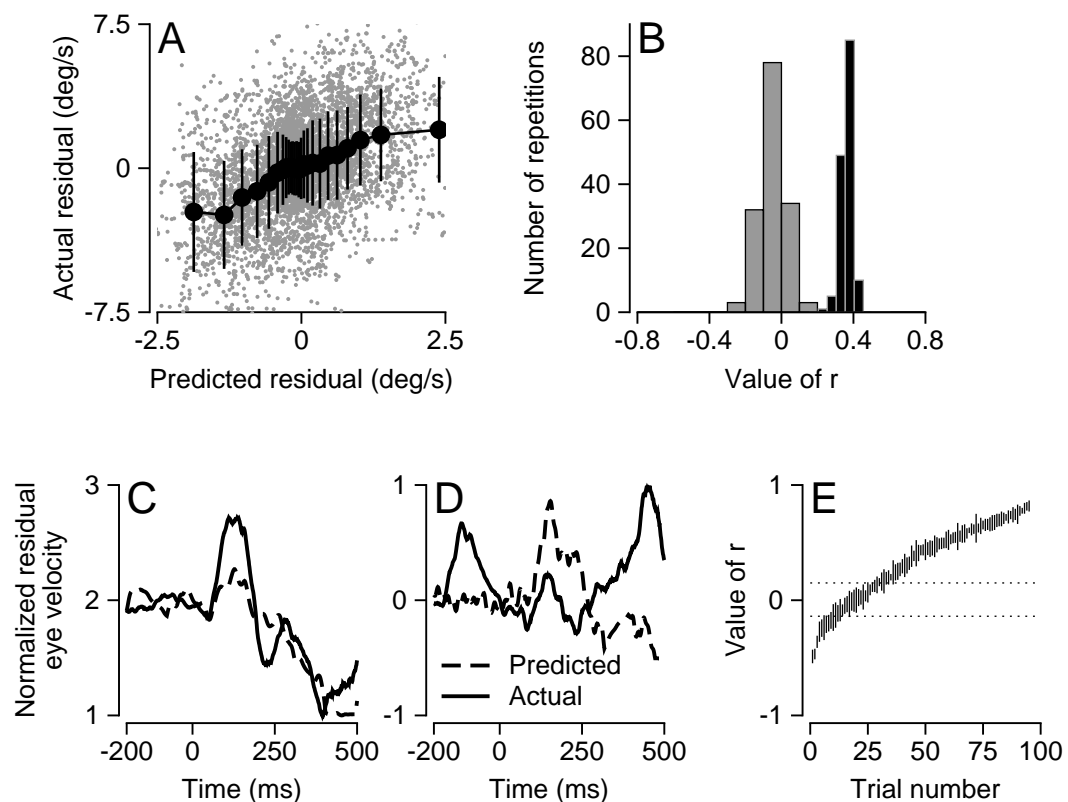


Figure 3.3: Predictions of the model relating spikes and behavior for the example neuron. A, Gray points show the actual and predicted residual eye velocity plotted for all time points across all trials. Black points represent the mean and standard deviation for the data shown in gray, grouped such that each of 20 bins would have an equal number of points. B, Histograms showing the distribution of correlations between actual and predicted residual eye velocity for 150 derivations of the filter based on different subsets of the data. Black and gray show results for the actual and shuffled data. C,D, Predicted residual eye velocity for the two trials yielding the best and worst correlations with actual eye velocity. Dashed and solid lines show the predicted and actual eye velocities. E, The distribution of correlations for each individual trial, ranked according to the value of r . The thickness of the ribbon shows one standard deviation above and below the mean correlation coefficient. The two horizontal dotted lines show the mean +2 standard deviations for the distribution of predictions made by the filter computed from shuffled data.

the outermost binned markers are likely compromised due to undersampling. For the full set of points, the correlation between the actual and predicted values of the residuals was 0.41. Plotting our data as in Supplementary Figure 3.3A supports our choice of a linear model for the relationship between residuals of neural responses and behavior.

To avoid overfitting, we used different sets of data to fit and test the model, across many runs. In Figures 2.2B and 2.2C, for example, the width of the black ribbon represents the mean and standard deviation of the filters derived from many draws of 40% (38 of 95) of our trials. The analysis described by Supplementary Figure 3.3A was repeated 150 times to provide statistically valid estimates of the performance of the linear filters. The distribution of correlations across runs was narrow (Supplementary Figure 3.3B, black bars) and was distinctly different from the distribution provided by analysis of shuffled data (Supplementary Figure 3.3B, gray bars). For the example neuron, the mean values of r across runs for the direct and shuffled analysis were 0.38 (SD: 0.03) and 0.0 (SD: 0.07). The holistic look at the correlation across the trials reassures us that our correlations between the predicted and actual values do not arise from the stereotyped nature of repeated movements.

Finally, we summarize the predictions of the filter on a trial-by-trial basis. For the example neuron, Supplementary Figures 3.3C and 3.3D show the actual (continuous line) and predicted (dashed line) residuals of eye velocity for the two trials with the largest ($r = 0.84$), and smallest ($r = -0.43$) correlation coefficients between actual and predicted residual eye velocity. Data has been normalized to the peak of the absolute value of each trace. Although the model does not capture all the details of the movement residuals, it does capture some aspects, particularly near the center of the trials. To obtain a single measure of the covariance between neural firing and behavior for each individual trial, we computed the correlations between model predictions and eye movement residuals for all trials (Supplementary Figure 3.3E), and compared these to correlations two standard deviations above predictions made from the shuffled filters (dashed lines). The majority of

trials ($n=60$) lie above the range of correlations for shuffled filters (dashed lines), and 8 lie below. The mean correlation coefficient was 0.34 (SD: 0.35). We conclude for the example neuron that the filter that relates the residual firing rate to the residual eye velocity has considerable explanatory power.

3.1.5 Generalization of neuron-behavior correlations to mean responses

For control analyses, we re-ran the filter-generation procedure on three different sets of data: i) responses to targets that moved 45 degrees above and below the preferred direction, ii) using the position trace, and iii) after shifting the data in time to minimize the mean squared difference between the eye velocity responses across trials. We also asked whether the filters obtained from the residuals of eye velocity and firing rate also could predict the mean eye movement. For the latter analysis, we convolved the filter obtained from the residuals for each neuron with the complete spike train and calculated the correlation coefficient with the actual eye movement. To evaluate filter performance when predicting the mean eye movement, we computed the spike-triggered average eye movement to serve as a benchmark:

$$Q_{ep}(t) = \frac{1}{n} \sum_{i=1}^n [e_i(t) \otimes \rho_i(t)] \quad (3.1)$$

We used this filter in conjunction with the complete spike train to generate a prediction for the complete eye movement for each trial as previously described. While this model performs well, it is constrained by the stereotyped nature of the behavior. Nonetheless, it provides an important reference point on how well linear filters might do at predicting behavior in a given direction [Machens et al., 2004]. Not surprisingly, the correlation coefficient between the filter predictions and the actual eye movement was always above zero, averaging 0.93 (SD:0.11). The filters we derived from the residuals constitute local, linear models of how the eye movement that is related to one extra spike in one neuron within

the limited space of trial-to-trial movement variations. To ask whether the filters generalize to the complete range of motions contained in a trial, we first examined how well they predicted the mean eye velocity from the mean spike train.

To generate a prediction of the mean eye velocity, we convolved the residual-derived filter with the mean firing rate (the PSTH in 1-ms bins). We then computed the correlation between this prediction and actual mean eye velocity as a function of time. As the filters were generated without respect to the mean, this method provides independent verification of our model. In general, statistically significant filters produced positive correlations (Supplementary Figure 3.4A, black bars), whereas non-significant filters produced a wide range of correlations with a distribution that surrounded zero (gray bars). The two distributions were different (Wilcoxon rank-sum, $p < 0.05$). Curiously, there were a handful of neurons with significant filters that yielded negative correlations with the mean eye movement but positive correlations with the residuals. We do not have a comprehensive explanation for this surprising, but infrequent, finding.

For each neuron, we compared the performance of the filter computed from the residuals with the best possible performance for a linear model by computing a performance index (PI):

$$PI = \frac{r_{\text{residuals}}}{r_{\text{STA}}} \quad (3.2)$$

where $r_{\text{residuals}}$ and r_{STA} are the correlations between the actual and predicted mean eye movements based on the filters obtained from the residuals versus from the spike-triggered average. Thus, PI describes the performance of the filter based on the residuals as a fraction of the maximum variance that can be captured by any linear filter. Values of PI close to 1 indicate that the predictive power of the filter derived from the residuals came closer to the best possible linear prediction. As before, a large portion of significant filters (Supplementary Figure 3.4B, black bars) performed significantly better than did non-significant filters (gray bars). Median values for significant versus non-significant filters were 0.32 (SD: 0.49)

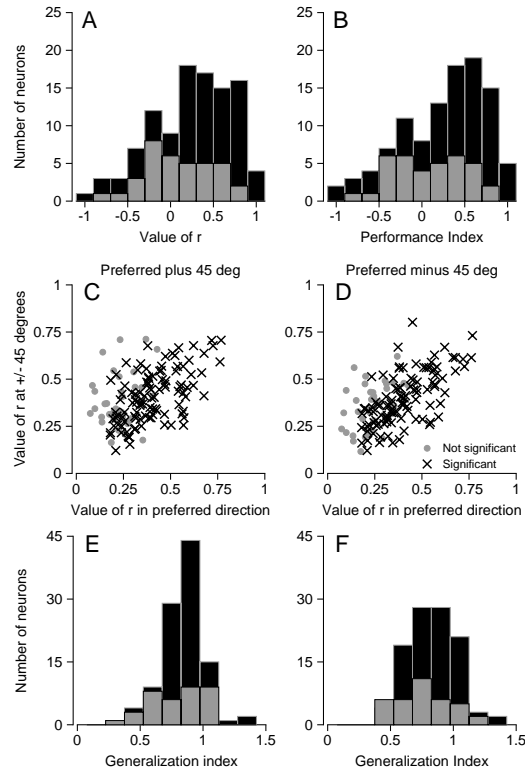


Figure 3.4: Generalization from filters derived for pursuit in the preferred direction to the mean eye movement and the eye movement variation in different directions. A, Distribution of the correlation coefficient between the prediction of the filter made on the mean eye movement and the actual mean eye movement. B, Distribution of the value of PI, which describes the correlation between actual and predicted mean eye velocity as a fraction of that for the optimal linear filter. C,D, The correlation between the predicted and actual eye velocity for pursuit in non-preferred directions is plotted versus that for pursuit in the preferred direction, where predictions were all based on filters derived from pursuit in the preferred direction. The direction of pursuit was 45 degrees clockwise (C) or counterclockwise (D) relative to the preferred direction. E,F, The distribution of the value of GI for the data in C and D. In all graphs, data from neurons with statistically significant filters are shown in black, and those from non-significant filters in gray.

versus 0.10 (SD: 0.42). The difference was statistically significant (Wilcoxon rank-sum test, $p < 0.05$). We conclude that a substantial subset of the filters derived from the residual eye movements describe the main features of the mean eye movement. As such, the filters derived from analysis of the residual spike trains and eye velocities are likely an excellent approximation of how the features in the eye movement are related to neural activity.

3.1.6 Generalization of neuron-behavior correlations across pursuit direction

Although we have analyzed so far only the data recorded during pursuit in the preferred direction of neurons in the FEF_{SEM}, we recorded responses during pursuit in eight directions spaced evenly around 360 deg. Neurons in the FEF_{SEM} tend to be broadly tuned, generally yielding responses vigorous enough to analyze in 3 of the 8 directions, including the preferred direction and the two directions rotated 45 deg clockwise and counter-clockwise from preferred. We found a strong relationship between the performance of filters on data from the preferred direction and the +45 deg directions.

For each neuron, we used the filters generated from residuals of spike trains and eye velocity in the neuron's preferred direction to compute predicted residuals of eye velocity for the +45 deg directions. The correlations between the actual and predicted residuals of eye velocity were similar across directions (Supplementary Figures 3.4C and 3.4D). To quantify how well the filters derived from the preferred direction generalized to other directions, we compared their performance to the performance of filters computed directly on the residuals of neural activity and eye velocity for each non-preferred direction. The two filters together allowed us to compute a generalization index (GI) for each neuron:

$$GI = \frac{r_{\text{pref}}}{r_{45}} \quad (3.3)$$

r_{pref} is performance of the filter derived in the preferred direction, evaluated on the non-preferred direction and r_{45} is the performance of a filter built and tested on the non-

preferred direction. A GI of one implies that the filter derived from the preferred direction data performed as well at describing activity in the non-preferred direction as a filter built directly from that data set. For pursuit in both the +45 and -45 deg directions (Supplementary Figures 3.4E and 3.4F), the distribution of GI was fairly close to one, indicating that the two filters made largely identical predictions, and by extension, that filters derived from pursuit in the preferred direction generalized well.

3.1.7 Relationship between filter performance and neural response properties

As with any population of cortical neurons, those selected for our study showed a range of response properties. However, we did not find evidence that the predictive value of the filter derived from residuals during pursuit in the preferred direction was related to any of the following response properties: the preferred direction (Supplementary Figure 3.5A), the tuning bandwidth (3.5B), the average firing rate during pursuit in the preferred direction (3.5C), the number of trials we were able to record for each direction of target motion (3.5D), the variance in eye velocity during pursuit (3.5E), and the variance in the spike train during pursuit (3.5F). We compared each response parameter to the strength of the relationship between spikes and eye velocity, as measured by the peak value of the correlation between the actual and predicted residuals of eye velocity, estimated in time. No graph showed a convincing relationship between the value of the measured parameter and either the strength of the neuron-behavior correlation or the likelihood of obtaining a statistically significant filter. Only the relationship between neuron-behavior correlation and tuning bandwidth was statistically significant: Figure 3.5B ($r = -0.39$, $p < 0.01$). For the other parameters: 3.5A (Mardia's $r = 0.02$, $p > 0.6$), 3.5C ($r = 0.10$, $p = 0.30$), 3.5D ($r = -0.04$, $p = 0.63$), 3.5E ($r = -0.02$, $p = 0.87$), 3.5F ($r = 0.16$, $p = 0.10$).

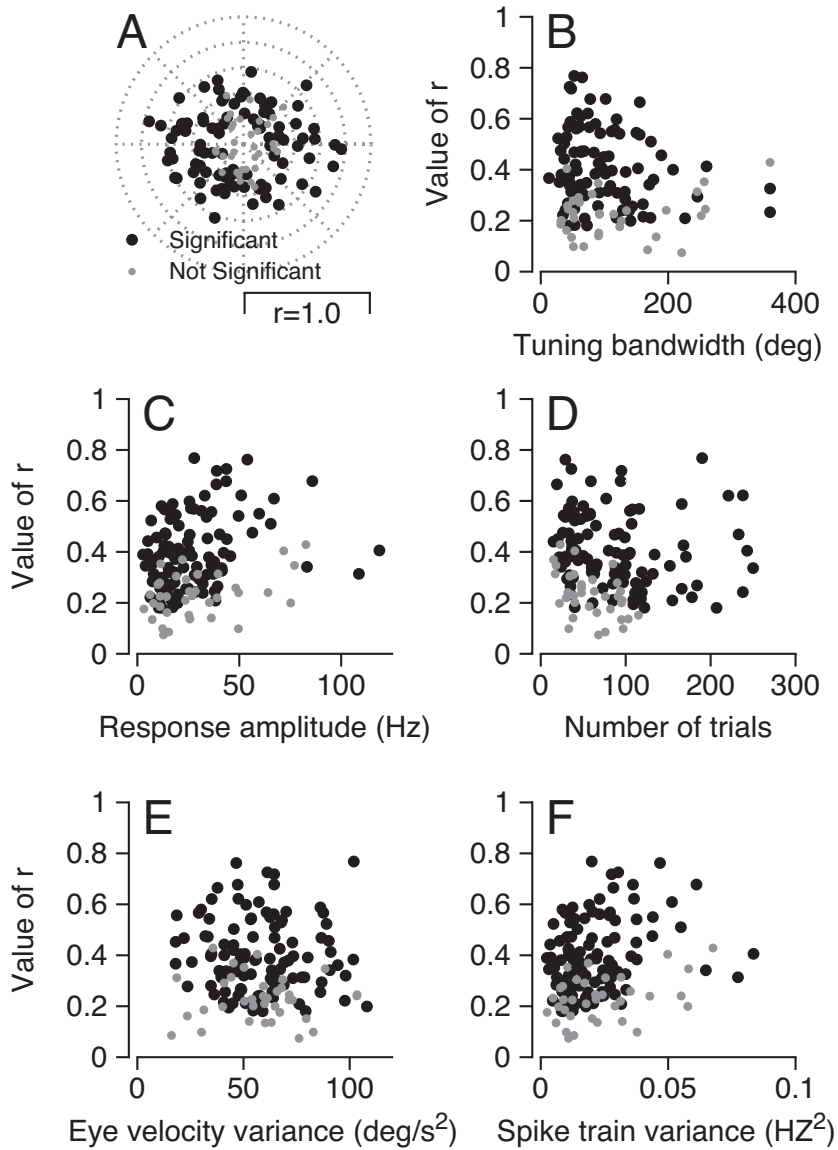


Figure 3.5: The relationship between traditional neural response properties and the correlation between actual and predicted residual eye velocity. Each graph plots correlation as a function of a different parameter of the neural response or experiment: A, preferred direction, B, tuning bandwidth, C, average response amplitude during pursuit, D, number of trials, E, total eye velocity variance, and F, total spike train variance.




Figure 3.6: Principal component analysis decomposition of the filters. A, Scree plot of eigenvalues, in black, scaled to report the fraction of variance described. The first three are color-coded green, blue, and red. B, The first three eigenvectors, each color corresponding to the plot in A. C, The projection of the filters into the space defined by the first two principal components. Filters have been color coded according to the original classification, where blue points are filters we called "single lobed," green "double," and red "triple."

3.2 Complementary Analyses

3.2.1 PCA decomposition of the filters

In order to confirm our categorization of our filters into single, double, and triple-lobed groups, we performed decomposed the unaligned matrix of filters into its principal components. We observed that the first three eigenvectors were single, double, and triple-lobed in shape, and together explained 39% of the variance. As these vectors matched our initial observations, we felt confident in using these three templates as archetypes. We then aligned our filters as described in the test, and re-ran the PCA decomposition. Supplementary Figure 3.6A shows a scree plot of eigenvalues, in black, scaled to report the fraction of variance described. The first three are color-coded green, blue, and red, to facilitate comparison with the corresponding eigenvectors, plotted in Supplementary Figure 3.6B.

To demonstrate the separation, Supplementary Figure 3.6C plots the projection of the filters into the space defined by the first two principal components. Points are color-coded to match our original classification (i.e. blue points are from filters we called single-lobed, green double, and red triple). The clear clustering confirms that our categorization was identical to an unbiased separation of filters.

3.2.2 Evaluation of cross-validation fractions for filter generation and testing

To avoid over-fitting, we split the data for each neuron into two fractions: one to generate a filter, and the other reserved for testing. Our final estimate of the strength of the correlation between the actual and predicted residuals of eye velocity depends on the fraction used: a filter based on too few trials is not an adequate representation of the data; conversely, a filter based on too high a fraction of the trials will be excellent, but without a sufficient test-bed. To find an appropriate compromise, we varied the percentage of trials devoted to generating filters, from 10 to 90% of the trials for each cell.

Supplementary Figure 3.78 shows that the choice to use 40% of the trials for filter derivation and 60% for test trials was an excellent compromise. Each gray symbol shows data for a single neuron and plots the variance of the correlation between predicted and actual residuals of eye velocity across multiple derivations of the filter as a function of the variance in the filter itself, integrated across its duration. Moving from left to right, the clusters of points and their means (black symbols and error bars) show results when 90% to 10% of the trials were used for filter generation and the remainder as test trials. Minimum variance of the neuron-behavior correlation falls at the point where 40% of the trials were used for filter generation, validating the approach taken in the rest of the paper.

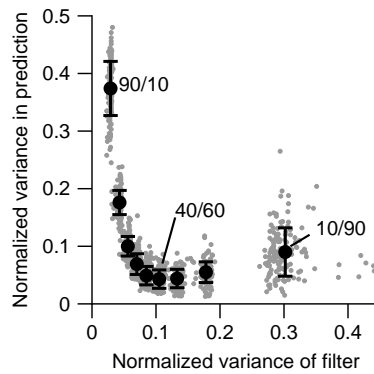


Figure 3.7: The effects on filter reliability of changing the ratio of trials used for filter generation to the fraction reserved for testing. Each point plots the normalized variance of the correlation between actual and predicted residuals of eye velocity versus the normalized variance of the filter itself. From left to right, the filters were generated on the basis of 90%, 80%, ..., 10% of the trials and tested on the remaining fraction. Data from each neuron is plotted nine times, once for each ratio. Gray symbols show the individual runs for each neuron and the open symbols with error bars show the mean and standard deviation across all neurons for a given ratio. Runs with 10%, 40%, and 90% of the trials used in filter generation are indicated by "90/10", "40/60" and "10/90".

Bibliography

- SA Adler, J Bala, and RJ Krauzlis. Primacy of spatial information in guiding target selection for pursuit and saccades. *J Vis*, 2(9):627–644, 2002.
- W Becker and R Jurgens. An analysis of the saccadic system by means of double step stimuli. *Vision Res*, 19(9):967–983, 1979.
- JW Bisley and ME Goldberg. Neural correlates of attention and distractibility in the lateral intraparietal area. *J Neurophysiol*, 95(3):1696–1717, 2006.
- M Brecht, M Schneider, B Sakmann, and T W Margrie. Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature*, 427(6976):704–710, Feb 2004. doi: 10.1038/nature02266. URL <http://www.hubmed.org/display.cgi?uids=14973477>.
- F Bremmer, W Graf, S Ben Hamed, and JR Duhamel. Eye position encoding in the macaque ventral intraparietal area (vip). *Neuroreport*, 10(4):873–878, 1999.
- Carlos D Brody. Disambiguating different covariation types. *Neural Comput*, 11(7):1527–1535, Oct 1999. JOURNAL ARTICLE.
- C J Bruce, M E Goldberg, M C Bushnell, and G B Stanton. Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J Neurophysiol*, 54(3):714–734, Sep 1985.

- CD Carello and RJ Krauzlis. Manipulating intent: evidence for a causal role of the superior colliculus in target selection. *Neuron*, 43(4):575–583, 2004.
- MR Carey, JF Medina, and SG Lisberger. Instructive signals for motor learning from visual cortical area mt. *Nat Neurosci*, 8(6):813–819, 2005.
- J M Carmena, M A Lebedev, C S Henriquez, and M A Nicolelis. Stable ensemble performance with single-neuron variability during reaching movements in primates. *J Neurosci*, 25(46):10712–10716, Nov 2005. doi: 10.1523/JNEUROSCI.2772-05.2005. URL <http://www.hubmed.org/display.cgi?uids=16291944>.
- E J Chichilnisky. A simple white noise analysis of neuronal light responses. *Network*, 12(2):199–213, May 2001.
- I-Han Chou and Stephen G Lisberger. The role of the frontal pursuit area in learning in smooth pursuit eye movements. *J Neurosci*, 24(17):4124–4133, Apr 2004.
- M M Churchland, B M Yu, S I Ryu, G Santhanam, and K V Shenoy. Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci*, 26(14):3697–3712, Apr 2006a. doi: 10.1523/JNEUROSCI.3762-05.2006. URL <http://www.hubmed.org/display.cgi?uids=16597724>.
- Mark M Churchland, Afsheen Afshar, and Krishna V Shenoy. A central source of movement variability. *Neuron*, 52(6):1085–1096, Dec 2006b.
- CE Connor, DC Preddie, JL Gallant, and DC Van Essen. Spatial attention effects in macaque area v4. *J Neurosci*, 17(9):3201–3214, 1997.
- EP Cook and JH Maunsell. Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. *J Neurosci*, 22(5):1994–2004, 2002.

- P Dayan and LF Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press, 2005.
- J J DiCarlo and J H Maunsell. Using neuronal latency to determine sensory-motor processing pathways in reaction time tasks. *J Neurophysiol*, 93(5):2974–2986, May 2005. doi: 10.1152/jn.00508.2004. URL <http://www.hubmed.org/display.cgi?uids=15548629>.
- R Dodge. An experimental study of visual fixation. *Psychological Review*, 8(4):1–88, 1907.
- JR Duhamel, CL Colby, and ME Goldberg. The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255(5040):90–92, 1992.
- JL Gardner and SG Lisberger. Linked target selection for saccadic and smooth pursuit eye movements. *J Neurosci*, 21(6):2075–2084, 2001.
- JL Gardner and SG Lisberger. Serial linkage of target selection for orienting and tracking eye movements. *Nat Neurosci*, 5(9):892–899, 2002.
- ME Goldberg and RH Wurtz. Activity of superior colliculus in behaving monkey. ii. effect of attention on neuronal responses. *J Neurophysiol*, 35(4):560–574, 1972.
- J P Gottlieb, C J Bruce, and M G MacAvoy. Smooth eye movements elicited by microstimulation in the primate frontal eye field. *J Neurophysiol*, 69(3):786–799, Mar 1993.
- Richard H R Hahnloser, Alexay A Kozhevnikov, and Michale S Fee. An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature*, 419(6902):65–70, Sep 2002.
- DM Halliday, JR Rosenberg, AM Amjad, P Breeze, BA Conway, and SF Farmer. A framework for the analysis of mixed time series/point process data—theory and application to the study of physiological tremor, single motor unit dis-

- charges and electromyograms. *Prog Biophys Mol Biol*, 64(2-3):237–278, 1995. URL <http://www.hubmed.org/display.cgi?uids=8987386>.
- C M Harris and D M Wolpert. Signal-dependent noise determines motor planning. *Nature*, 394(6695):780–784, Aug 1998.
- E G Keating. Frontal eye field lesions impair predictive and visually-guided pursuit eye movements. *Exp Brain Res*, 86(2):311–323, 1991.
- RJ Krauzlis and SG Lisberger. Temporal properties of visual motion signals for the initiation of smooth pursuit eye movements in monkeys. *J Neurophysiol*, 72(1):150–162, 1994.
- D Lee, N L Port, W Kruse, and A P Georgopoulos. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci*, 18(3):1161–1170, Feb 1998.
- SG Lisberger. Postsaccadic enhancement of initiation of smooth pursuit eye movements in monkeys. *J Neurophysiol*, 79(4):1918–1930, 1998.
- SG Lisberger and LE Westbrook. Properties of visual inputs that initiate horizontal smooth pursuit eye movements in monkeys. *J Neurosci*, 5(6):1662–1673, 1985.
- SG Lisberger, AF Fuchs, WM King, and LC Evinger. Effect of mean reaction time on saccadic responses to two-step stimuli with horizontal and vertical components. *Vision Res*, 15:1021–1025, 1975.
- D Liston and RJ Krauzlis. Shared response preparation for pursuit and saccadic eye movements. *J Neurosci*, 23(36):11305–11314, 2003.
- D Liston and RJ Krauzlis. Shared decision signal explains performance and timing of pursuit and saccadic eye movements. *J Vis*, 5(9):678–689, 2005.
- AE Luebke and DA Robinson. Transition dynamics between pursuit and fixation suggest different systems. *Vision Res*, 28(8):941–946, 1988.

- Rogelio Luna, Adrian Hernandez, Carlos D Brody, and Ranulfo Romo. Neural codes for perceptual discrimination in primary somatosensory cortex. *Nat Neurosci*, 8(9):1210–1219, Sep 2005. Comparative Study.
- Wei Ji Ma, Jeffrey M Beck, Peter E Latham, and Alexandre Pouget. Bayesian inference with probabilistic population codes. *Nat Neurosci*, 9(11):1432–1438, Nov 2006.
- Christian K Machens, Michael S Wehr, and Anthony M Zador. Linearity of cortical receptive fields measured with natural sounds. *J Neurosci*, 24(5):1089–1100, Feb 2004.
- Z F Mainen and T J Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, Jun 1995. In Vitro.
- A Maye, C H Hsieh, G Sugihara, and B Brembs. Order in spontaneous behavior. *PLoS ONE*, 2, 2007. doi: 10.1371/journal.pone.0000443. URL <http://www.hubmed.org/display.cgi?uids=17505542>.
- E M Maynard, N G Hatsopoulos, C L Ojakangas, B D Acuna, J N Sanes, R A Normann, and J P Donoghue. Neuronal interactions improve cortical population coding of movement direction. *J Neurosci*, 19(18):8083–8093, Sep 1999.
- LE Mays and DL Sparks. Saccades are spatially, not retinocentrically, coded. *Science*, 208(4448):1163–1165, 1980.
- JF Medina, MR Carey, and SG Lisberger. The representation of time for motor learning. *Neuron*, 45(1):157–167, Jan 2005.
- T Moore and KM Armstrong. Selective gating of visual signals by microstimulation of frontal cortex. *Nature*, 421(6921):370–373, 2003.
- MC Morrone, J Ross, and D Burr. Saccadic eye movements cause compression of time as well as space. *Nat Neurosci*, 8(7):950–954, 2005.

- VB Mountcastle, RA Andersen, and BC Motter. The influence of attentive fixation upon the excitability of the light-sensitive neurons of the posterior parietal cortex. *J Neurosci*, 1(11):1218–1225, 1981.
- JR Muller, MG Philiastides, and WT Newsome. Microstimulation of the superior colliculus focuses attention without moving the eyes. *Proc Natl Acad Sci U S A*, 102(3):524–529, 2005.
- KI Nagel and AJ Doupe. Temporal processing and adaptation in the songbird auditory forebrain. *Neuron*, 51(6):845–859, Sep 2006.
- WT Newsome, RH Wurtz, MR Dursteler, and A Mikami. Deficits in visual motion processing following ibotenic acid lesions of the middle temporal visual area of the macaque monkey. *J Neurosci*, 5(3):825–840, 1985.
- WT Newsome, RH Wurtz, and H Komatsu. Relation of cortical areas mt and mst to pursuit eye movements. ii. differentiation of retinal from extraretinal inputs. *J Neurophysiol*, 60(2):604–620, 1988.
- T Ogawa and M Fujita. Velocity profile of smooth pursuit eye movements in humans: pursuit velocity increase linked with the initial saccade occurrence. *Neurosci Res*, 31(3):201–209, 1998.
- L C Osborne, S G Lisberger, and W Bialek. A sensory source for motor variation. *Nature*, 437(7057):412–416, Sep 2005. doi: 10.1038/nature03961. URL <http://www.hubmed.org/display.cgi?uids=16163357>.
- Leslie C Osborne, William Bialek, and Stephen G Lisberger. Time course of information about motion direction in visual area MT of macaque monkeys. *J Neurosci*, 24(13):3210–3222, Mar 2004.

- Leslie C Osborne, Sonja S Hohl, William Bialek, and Stephen G Lisberger. Time course of precision in smooth-pursuit eye movements of monkeys. *J Neurosci*, 27(11):2987–2998, Mar 2007. Comparative Study.
- Liam Paninski, Matthew R Fellows, Nicholas G Hatsopoulos, and John P Donoghue. Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *J Neurophysiol*, 91(1):515–532, Jan 2004a. Comparative Study.
- Liam Paninski, Shy Shoham, Matthew R Fellows, Nicholas G Hatsopoulos, and John P Donoghue. Superlinear population encoding of dynamic hand trajectory in primary motor cortex. *J Neurosci*, 24(39):8551–8561, Sep 2004b.
- K G Pearson. Common principles of motor control in vertebrates and invertebrates. *Annu Rev Neurosci*, 16:265–297, 1993. doi: 10.1146/annurev.ne.16.030193.001405. URL <http://www.hubmed.org/display.cgi?uids=8460894>.
- Jonathan W Pillow and Eero P Simoncelli. Dimensionality reduction in neural models: an information-theoretic generalization of spike-triggered average and covariance analysis. *J Vis*, 6(4):414–428, 2006.
- A Pouget and TJ Sejnowski. A neural model of the cortical representation of egocentric distance. *Cereb Cortex*, 4(3):314–329, 1994.
- P Pouget, E E Emeric, V Stuphorn, K Reis, and J D Schall. Chronometry of visual responses in frontal eye field, supplementary eye field, and anterior cingulate cortex. *J Neurophysiol*, 94(3):2086–2092, Sep 2005. doi: 10.1152/jn.01097.2004. URL <http://www.hubmed.org/display.cgi?uids=15944228>.
- R Ramachandran and SG Lisberger. Normal performance and expression of learning in the vestibulo-ocular reflex (vor) at high frequencies. *J Neurophysiol*, 93(4):2028–2038, 2005.

- JB Reppas, WM Usrey, and RC Reid. Saccadic eye movements modulate visual responses in the lateral geniculate nucleus. *Neuron*, 35(5):961–974, 2002.
- A Riehle and J Requin. The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex. *Behav Brain Res*, 53(1-2): 35–49, Feb 1993.
- Fred Rieke, David Warland, R R de Ruyter van Steveninck, and William Bialek. *Spikes: Exploring the Neural Code*. The MIT Press, 1999.
- D L Ringach, G Sapiro, and R Shapley. A subspace reverse-correlation technique for the study of visual neurons. *Vision Res*, 37(17):2455–2464, Sep 1997. URL <http://www.hubmed.org/display.cgi?uids=9381680>.
- Nicole C Rust, Valerio Mante, Eero P Simoncelli, and J Anthony Movshon. How MT cells analyze the motion of visual patterns. *Nat Neurosci*, 9(11):1421–1431, Nov 2006.
- B Sakitt. Counting every quantum. *J Physiol*, 223(1):131–150, May 1972.
- E Salinas and LF Abbott. Coordinate transformations in the visual system: how to generate gain fields and what to compute with them. *Prog Brain Res*, 130:175–190, 2001.
- A Schlack, KP Hoffmann, and F Bremmer. Selectivity of macaque ventral intraparietal area (area vip) for smooth pursuit eye movements. *J Physiol*, 551(Pt 2):551–561, 2003.
- David Schoppik and Stephen G Lisberger. Saccades exert spatial control of motion processing for smooth pursuit eye movements. *J Neurosci*, 26(29):7607–7618, Jul 2006.
- M N Shadlen and W T Newsome. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*, 18(10):3870–3896, May 1998. URL <http://www.hubmed.org/display.cgi?uids=9570816>.

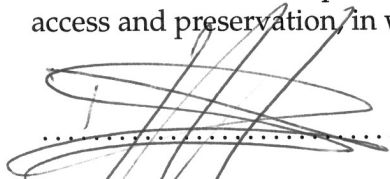
- M N Shadlen, K H Britten, W T Newsome, and J A Movshon. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci*, 16(4):1486–1510, Feb 1996.
- Tatyana Sharpee, Nicole C Rust, and William Bialek. Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput*, 16(2):223–250, Feb 2004.
- MA Smith and JD Crawford. Distributed population mechanism for the 3-d oculomotor reference frame transformation. *J Neurophysiol*, 93(3):1742–1761, 2005.
- JF Soechting, LA Mrotek, and M Flanders. Smooth pursuit tracking of an abrupt change in target direction: vector superposition of discrete responses. *Exp Brain Res*, 160(2):245–258, 2005.
- M Tanaka and SG Lisberger. Enhancement of multiple components of pursuit eye movement by microstimulation in the arcuate frontal pursuit area in monkeys. *J Neurophysiol*, 87(2):802–818, 2002a.
- Masaki Tanaka and Stephen G Lisberger. Role of arcuate frontal cortex of monkeys in smooth pursuit eye movements. I. Basic response properties to retinal image motion and position. *J Neurophysiol*, 87(6):2684–2699, Jun 2002b.
- F E Theunissen. From synchrony to sparseness. *Trends Neurosci*, 26(2):61–64, Feb 2003. URL <http://www.hubmed.org/display.cgi?uids=12536128>.
- F E Theunissen, K Sen, and A J Doupe. Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *J Neurosci*, 20(6):2315–2331, Mar 2000.
- Emanuel Todorov and Michael I Jordan. Optimal feedback control as a theory of motor coordination. *Nat Neurosci*, 5(11):1226–1235, Nov 2002.
- D J Tolhurst, J A Movshon, and A F Dean. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res*, 23(8):775–785, 1983.

- AS Tolias, T Moore, SM Smirnakis, EJ Tehovnik, AG Siapas, and PH Schiller. Eye movements modulate visual receptive fields of v4 neurons. *Neuron*, 29(3):757–767, 2001.
- V J Uzzell and E J Chichilnisky. Precision of spike trains in primate retinal ganglion cells. *J Neurophysiol*, 92(2):780–789, Aug 2004.
- R VanRullen. A simple translation in cortical log-coordinates may account for the pattern of saccadic localization errors. *Biol Cybernetics*, 91:131–137, 2004.
- W E Vinje and J L Gallant. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287(5456):1273–1276, Feb 2000.
- Michael C-K Wu, Stephen V David, and Jack L Gallant. Complete functional characterization of sensory neurons by system identification. *Annu Rev Neurosci*, 29:477–505, 2006.
- Robert H Wurtz and Marc A Sommer. Identifying corollary discharges for movement in the primate brain. *Prog Brain Res*, 144:47–60, 2004.
- D Zipser and RA Andersen. A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331(6158):679–684, 1988.
- E Zohary, M N Shadlen, and W T Newsome. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370(6485):140–143, Jul 1994.

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses and dissertations. Copies of all UCSF theses and dissertations will be routed to the library via the Graduate Division. The library will make all theses and dissertations accessible to the public and will preserve these to the best of their abilities, in perpetuity.

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis or dissertation to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.



.....

Author Signature

6/21/07
.....

Date