DYNAMICS OF SENSORIMOTOR INTEGRATION IN GAZE CONTROL

by

Uday K. Jagadisan

Bachelor of Technology in Biotechnology,
Indian Institute of Technology Madras, 2006

Master of Engineering in Bioengineering,
Syracuse University, 2008

Submitted to the Graduate Faculty of

Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2017

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Uday K. Jagadisan

It was defended on

February 23, 2016

and approved by

Aaron Batista, Ph.D., Professor, Department of Bioengineering

Carl Olson, Ph.D., Professor, Center for the Neural Basis of Cognition
Carnegie Mellon University

Andrew Schwartz, Ph.D., Professor, Department of Neurobiology

Dissertation Director: Neeraj Gandhi, Ph.D., Professor, Department of Bioengineering

Copyright © by Uday K. Jagadisan

2017

DYNAMICS OF SENSORIMOTOR INTEGRATION IN GAZE CONTROL

Uday K. Jagadisan, Ph.D.

University of Pittsburgh, 2017

We are immersed in a rich and complex visual environment. Foveal animals such as primates interact with this environment primarily by shifting their gaze to objects of interest, an act colloquially known as "looking". Gaze shifts involve a sensory-to-motor transformation of the visual world into a movement command that redirects the line of sight. Several brain regions involved in the gaze control network enable this transformation by representing sensory, cognitive, and movement-related information in the same population of neurons. For example, neurons in the intermediate layers of the superior colliculus, a critical gaze control node in the midbrain, fire action potentials both in response to the onset of a visual target, during cognitive processing such as target selection and decision-making, and during an eye movement to that target. However, it is not known how this multiplexed dynamics of sensorimotor integration is read out by a downstream decoder. From the perspective of the decoder, when does sensory and cognitive processing end and movement programming begin?

We asked whether sensorimotor integration occurs in sequential stages or is implemented in parallel. Our approach was to 1) study the role of gaze fixation in shaping sensorimotor activity, 2) chart the time course of movement preparation by removing inhibitory gating on the network, and, 3) look for signatures of aforementioned multiplexing in neuronal population dynamics. Influential theories have proposed that sensory and movement processing occur in a cascade implemented by different subsets of neurons, and neural activity must accumulate to a pre-determined level in order to actuate the gaze shift. In contrast, we found that sensorimotor information is represented in a flexible continuum of neurons. Furthermore, ongoing low-frequency activity during the transformation has motor potential - latent ability to produce an eye

movement - that is revealed under the right conditions. Together, these results suggest that the brain performs sensorimotor integration in parallel. Finally, the temporal structure of population activity evolved differently for sensory and movement-related processes, providing a novel substrate by which downstream neurons may discriminate between the two in order to decide when to generate a movement.

TABLE OF CONTENTS

PREFACE ..... **XII**

ACKNOWLEDGEMENTS ..... **XIII**

1.0 GENERAL INTRODUCTION – IMPORTANCE OF GAZE CONTROL **1**

1.1 GAZE CONTROL NETWORK IN THE PRIMATE BRAIN **1**

1.1.1 The Superior Colliculus **2**

1.2 SENSORIMOTOR INTEGRATION IN GAZE CONTROL **3**

1.2.1 Fixation and Saccade Generation **4**

1.2.2 Movement Preparation and Initiation **5**

1.3 RESEARCH OBJECTIVES **7**

1.3.1 Knowledge Gap **7**

1.3.2 Research Strategy **8**

2.0 DISRUPTION OF FIXATION REVEALS LATENT SENSORIMOTOR PROCESSES IN THE SUPERIOR COLLICULUS **9**

2.1 INTRODUCTION **9**

2.2 METHODS AND DATA ANALYSIS **11**

2.2.1 General and surgical procedures **11**

2.2.2 Visual stimuli and behavior **11**

2.2.3 Perturbations **12**

2.2.4 Electrophysiology **12**

2.2.5 Data analysis and pre-processing **13**

2.3 RESULTS 15

2.3.1 Sub-threshold enhancement of stimulus-evoked activity following disruption of fixation 17

2.3.2 A hidden visual response is present in putative movement neurons 19

2.3.3 Fixation-related activity in rostral SC is dissociated from visuomovement activity 20

2.3.4 Rostral SC activity is better correlated with microsaccade occurrence ... 22

2.3.5 Alternate perturbations reveal mechanism mediating fixation-saccade network balance 24

2.3.6 Activity during premature saccades is reflective of a caudal SC-driven shift in balance 27

2.4 DISCUSSION 30

3.0 EXTRACTING SACCADES FROM EYE MOVEMENTS TRIGGERED BY REFLEX BLINKS 34

3.1 INTRODUCTION 34

3.2 METHODS 37

3.2.1 General and surgical procedures 37

3.2.2 Visual stimuli and behavior 37

3.2.3 Induction of reflex blinks 38

3.2.4 Electrophysiology 38

3.3 DATA ANALYSIS 39

3.3.1 Data analysis and pre-processing 39

3.3.2 Modeling of blink-triggered movements 39

3.3.3 Detection of back-thresholded saccades 40

3.3.4 Other analyses and statistical tests 41

3.4 RESULTS 42

3.5 DISCUSSION 50

4.0 REMOVAL OF INHIBITION UNCOVERS LATENT SACCADE PREPARATION DYNAMICS 53

4.1 INTRODUCTION 53

4.2 METHODS AND DATA ANALYSIS 56

4.2.1 General and surgical procedures 56

4.2.2 Visual stimuli and behavior 56

4.2.3 Induction of reflex blinks 57

4.2.4 Electrophysiology 57

4.2.5 Data analysis and pre-processing 58

4.2.6 Linear decomposition of blink-triggered movements 58

4.2.7 Surrogate data analysis 59

4.3 RESULTS 59

4.3.1 The blink perturbation triggers reduced latency saccades 60

4.3.2 Blink-triggered saccades are evoked at lower thresholds compared to normal saccades 61

4.3.3 Rate of accumulation of SC activity accelerates following disinhibition by the blink 63

4.3.4 SC preparatory activity preceding the saccade-related burst possesses motor potential 65

4.4 DISCUSSION 67

4.4.1 Implications for threshold-based accumulator models 67

4.4.2 Alternative models of movement initiation 68

4.4.3 Parallel implementation of the sensory-to-motor transformation 69

5.0 POPULATION TEMPORAL STRUCTURE SUPPLEMENTS THE RATE CODE DURING SENSORIMOTOR TRANSFORMATIONS 71

5.1 ABSTRACT 71

5.2 INTRODUCTION 72

5.3 METHODS 74

5.3.1 General and surgical procedures 74

5.3.2 Visual stimuli and behavior 74

5.3.3 Electrophysiology 75

5.4 DATA ANALYSIS 76

5.4.1 Preliminary analyses 76

5.4.2 Inferring population dynamics from single-unit recordings 77

5.4.3 Temporal stability analyses 78

5.4.4 Leaky accumulator with facilitation (LAF) model 80

5.5 RESULTS 82

5.6 DISCUSSION 88

6.0 DISCUSSION AND CONCLUSIONS 90

6.1 THE SENSORIMOTOR MULTIPLEXING PROBLEM 90

6.2 POSSIBLE SOLUTIONS 91

6.3 SENSORIMOTOR MULTIPLEXING AND THE PREMOTOR THEORY OF ATTENTION 95

6.4 CONCLUDING REMARKS 97

APPENDIX A 98

BIBLIOGRAPHY 107

LIST OF FIGURES

Figure 21. Behavior in control and perturbation conditions 16

Figure 22. Population activity in caudal SC 18

Figure 23. Latent visual response in putative movement neurons 20

Figure 24. Population activity in rostral SC 21

Figure 25. Microsaccade behavior of individual animals 23

Figure 26. Behavior and population activity for alternative perturbations 25

Figure 27. Microsaccade behaviour during alternative perturbations 26

Figure 28. Population activity during premature saccades 29

Figure 31. Blink-related eye movements and blink-triggered gaze shifts 36

Figure 32. Saccade extraction algorithms 41

Figure 33. Comparison of model performance – behaviour 43

Figure 34. Robustness of model performance to saccade direction 45

Figure 35. Comparison of model performance – neural activity 46

Figure 36. Time course of extracted saccades and activity suppression 49

Figure 41. Testing models of movement preparation 55

Figure 42. Time course of blink-triggered saccades 60

Figure 43. Analysis of putative threshold 62

Figure 44. Analysis of accumulation rate change following perturbation 64

Figure 45. Correlation between neural activity and pre-saccadic velocity 66

Figure 51. Sensorimotor transformations are mediated by neurons that multiplex sensory and motor information 73

Figure 52. Population activity is temporally unstable in the visual burst and stable in the motor burst..... 83

Figure 53. The stability hypothesis is validated for other neural populations and tasks 85

Figure 54. Leaky accumulator with facilitation model can discriminate population temporal structure..... 87

Figure 61. Summary of models of movement preparation 94

Figure A1. Inferring population dynamics from single-unit recordings..... 99

Figure A2. Temporal stability as a function of $\boldsymbol{\tau}$ 100

Figure A3. Temporal stability profile is specific to the recorded population structure..... 101

Figure A4. Population activity of rostral SC neurons is stable during microsaccades 102

Figure A5. Characteristics of the population in ipsilateral SC 103

Figure A6. Temporal stability of combined populations..... 104

PREFACE

The chapters of this dissertation are standalone manuscripts that have either been published in some form, under review, or under preparation to be submitted.

General Introduction: Large sections have been published in the form of a book chapter.
Jagadisan, U.K. and Gandhi, N.J. "Neural mechanisms of target selection in the superior colliculus", In: The New Visual Neurosciences, edited by L.M. Chalupa and J.S. Werner, MIT Press, 2014.

Chapter 2: Jagadisan, U.K. and Gandhi, N.J. (2016) "Disruption of fixation reveals latent sensorimotor processes in the superior colliculus", Journal of Neuroscience, 36(22): 6129-40.

Chapter 3: Jagadisan, U.K. and Gandhi, N.J. "Extracting saccades from eye movements triggered by reflex blinks", Journal of Neurophysiology, 2016 [Under revision].

Chapter 4: Jagadisan, U.K. and Gandhi, N.J. "Removal of inhibition uncovers latent saccade preparation dynamics", 2016 [Under review].

Chapter 5: Jagadisan, U.K. and Gandhi, N.J. "Population temporal structure supplements the rate code during sensorimotor transformations", 2016 [Under revision].

Discussion: Manuscript under preparation for submission as a review/opinion article.

ACKNOWLEDGEMENTS

It is almost sacrilege to complete a monumentous project like a doctoral dissertation and not let the universe that conspired to make it happen know how grateful one is. Hence, what follow are sacred words.

As a student of the idea that most things in the universe are chaotic dynamical systems, with as yet unknown forces keeping things seemingly ordered, I would like to begin by acknowledging the most significant influences I have had in dictating the course of my life - my initial conditions. It goes without saying that the most important people in my life are my family. My parents and my sister have been my everything for as long as I can remember, and their unwavering support through thick and thin has been a critical factor in me keeping my sanity through the course of my doctoral career. My father, Jagadisan, has been a source of moral strength, pragmatic guidance and spiritual inspiration (as much as I have veered from the "path"). A brilliant, hands on engineer (of which I am just a fraction), I owe my intellectual and analytical abilities to him. He is the quintessential tinkerer, and it is a surprise I am not in a harder engineering field, having grown up surrounded by tools and gadgets and geek literature of all sorts. Of course, we now know that there is nothing harder than the hard problem of the mind. My mother, Nirmala, has been a source of unconditional love, warmth and emotional guidance. A physics teacher, she instilled in me a sense of wonder of the world and all things physics. She taught me equanimity, the ability to not let emotions get the better of you under any circumstance, by example – an invaluable lesson to anyone who ever pursues a doctoral degree. My younger sister, Archana, has been my source of joy forever. I cannot imagine having had a childhood without her, and she has been my closest companion and best friend for a long time. Far more mature than I ever was and able to adopt the mantle of elder sibling at will, she has been a significant factor in keeping me on track. I cannot reiterate how invaluable their presence in my life is. Without their constant efforts, guidance, and goading, it is a fact of physical law that I would not be where I am today (wherever

this is). Finally, I would like to dedicate this work to my grandfather (thatha) – the man who coddled me and taught me values - who was eagerly waiting to call me "Dr.", but unfortunately passed away less than a month ago.

I would like to thank Dr.Neeraj (Raj) Gandhi, my graduate dissertation advisor (and academic father). A wonderful scientist, the work in this thesis was liberally inspired by his ideas. I have learnt the importance (and philosophical necessity) of treating all theories fairly and without preconceived notions or prejudice, from him. Raj is the perfect mentor, always there when you need him (just a knock on the door away), willing to have long chats about any unripe thought the young graduate student in me had, providing the long leash necessary for creative and scientific freedom, and controlling it when things started going astray, as was common with my bouts of indiscipline. Apart from scientific advice in the lab, Raj takes a personal interest in his students and ensures they are on the right track career-wise, which is not something that can be said for all advisors, based on the experience of my friends and colleagues. This dissertation would not exist without his guidance and support. I am honoured to extend my stint in his lab for a short while as a postdoc.

I would like to thank the members of my dissertation committee, Dr. Aaron Batista, Dr. Carl Olson, and Dr. Andy Schwartz, for serving on my committee and providing me with valuable guidance during committee meetings. It should be mentioned that Aaron, in particular, took a personal interest in me and offered to help me improve my presentation skills (which I have not taken up yet). Carl's incisive comments on my proposal and work-in-progress presentations at CNBC Brain Group were instrumental in shaping my thought process. Their comments during preparation of one of the manuscripts that contributed to this thesis went a long way in improving its quality. Andy's presence and work at Pitt was one of the initial attractions of attending graduate school here; his mentorship and camaraderie in a class I took with him has impacted my teaching philosophy.

A lot of the work in this thesis would not have been possible without the help and support of lab members, former and present. I would especially like to thank Gloria Foster, our lab manager for the first few years of my grad school stint, for her help in teaching me the ropes of daily life in a non-human primate lab, from initial acclimatization of (and to) monkeys, the chair training setup, and behavioural training to help with chamber cleanings, palliative care of animals, and lab management responsibilities. Her help during those years when the lab was just the two of us (later

dwindling to just me) and Raj, was invaluable in keeping things running smoothly. Husam, the first PhD student to graduate from the lab, was a huge help in teaching me the basics of single-unit electrophysiology, and the fact that he had set up the experimental rig into a usable shape made it easier for me to get started. The newer members of the lab – Ivan, Alex, Kevin, Michelle, and Corentin, helped make the last couple of years a really fun and enjoyable experience with their presence and camaraderie. Ivan has been a great friend and a budding neuroscientist with wonderful tinkering skills, and I have no doubt he will go on to become a successful researcher in his own right. Kevin's motivational pep talks helped me plod through to finish last-minute on site posters more than once – he has immense potential as a future neuroscientist as well. I would like to thank all of them for picking up the slack in lab management during my dissertation-writing months, and spending hours listening to my practice defense talks and helping me smoothen out the rough edges.

I would like to thank the administrative personnel in the Department of Bioengineering for their behind-the-scenes contributions to the completion of this degree. I would especially like to mention Lynette Spataro, Nick Mance (who succeeded Lynette as BioE Graduate Coordinator), and Glenn Peterson, for their numerous patient responses to my questions about departmental regulations, and impromptu reminders about timelines. I've probably repeatedly (and regretfully) made some of their lives harder by not sticking to deadlines or filing paperwork appropriately. I would also like to thank Dr. Harvey Borovetz, who was the departmental chair for the first half of my graduate career, and Dr. Sanjeev Shroff, the chair for the latter half. I gathered a lot of ideas about teaching style and philosophy from Sanjeev having TA-ed his course on Dynamical Systems in Physiology twice.

I would like to thank all the teachers and educators I have been fortunate to have learned from in the past, including, but not restricted to, my teachers in school (Varghese sir, my maths teacher in high school gets a special mention), the tutors at Ramaiah JEE coaching centre, and the professors and lecturers at IIT Madras. Prof. Vidyanand Nanjundiah, with whom I pursued my summer research project at IISc, was a great inspiration as a scientist and a thinker, and one of the first examples I encountered of people with a physics/quantitative background applying their knowhow in the messy domain of biology.

Although I was deeply interested in how the biology behind the brain comes to define us for as long as I've known about the brain, it was my work at Syracuse University modelling

episodic memory and semantic learning which created the realization that we need to know more about the system through experiments before (or in parallel with) attempting to simplify using theory. I would like to thank Dr. Marc Howard, my advisor at Syracuse, and Dr. Karthik Shankar, a friend and postdoc in the lab with whom I worked closely. Karthik was a great bouncing board for ideas during the early years of my doctoral career, and I cherish the deep discussions I've had with him about various aspects of neuroscience, physics, and the philosophy of science. I'm fortunate to have had the first two publications under my belt co-authored with Marc and Karthik.

I would like to thank my friends who've kept me company, and guided and inspired me to reach greater heights throughout my life. This includes, in reverse chronological order, all the friends I've known in Pittsburgh (far too many to individually list, but I owe special mentions to Krishna, fellow neuroscientist and my roommate for 3 years, and Samrat, a great buddy), including the quiz club gang and the Sanskrit/Ashrama gang, friends from Syracuse, friends from IIT (the BT gumbal, Dheeru), the IAS summer gang (esp., Nisha, who always believed in me), and school friends (Karthik, Chakrapani, Ananya, Aditya, Ankur). My time off from research with the Pittsburgh folks, mostly at quiz club or dinner/late night gatherings, were much-needed rejuvenation sessions.

I would like to reiterate that the motivation and moral support of my immediate and substantially large extended family, starting during childhood and all the way until the completion of this PhD, was crucial in helping me maintain my composure and balance during this arduous process. I would like to thank my parents and sister again, along with my grandparents (thatha and paatis), uncles (chitappas, perippas and mamas), aunts (chittis, perimas, and mamis), and numerous cousins for sharing my small accomplishment joys and failure sorrows along the way. I would like to specifically mention Shriram chitaps, who always takes pride in my accomplishments and is a constant source of encouragement, and keeps motivating me to reach for the stars. I would also like to share my joy with Akshaya, my young cousin sister, who is a budding biomedical researcher, and explaining neuroscience to whom has always been fun. Special shout out, amongst family friends, to Saraswati aunty, Venu uncle and family (whose house was my second home growing up), and CSK uncle.

Plenty of gratitude is due to my closest friends, classmates from IIT, from whom I'm only a sporadic phone call or phone tag sequence away to this day. Suri, my call-pal and fellow academic, is a source of inspiration for his intense focus and dedication to his research. He's been

like a brother for the past twelve years, and we've spent many a much-needed vacation away from the lab with each other, discussing life and relationships, sipping coffee and exploring the outdoors. SK, my fellow long haul grad student, and I have undergone significant personal evolution over the course of our graduate school years, and one of the parts of those years I will miss the most is the long phone conversations with him about life, science, and academia, and any and every meta-level thought stream emanating from there. Balaji, a fellow aspiring neuroscientist and I have had numerous stimulating conversations about systems neuroscience, and the paths forward to a professional career in the field. I'm glad that we get to rendezvous annually at neuroscience meetings. My chats with Vikash, my brother-in-arms, once aspiring cognitive scientist, and fellow romantic, psychonaut, armchair philosopher and questioner-of-things, have significantly impacted my thought process about the big questions we're all interested in, including the philosophy of neuroscience research. The impromptu panel discussions/dorm room sessions I shared with these folks about every imaginable topic under the sun, from science and engineering to information, randomness, intelligence, evolution, GoD, and the nature of reality itself (basically, life, the universe, and everything) have shaped our collective understanding and worldviews, with a part of that subtly but surely permeating into how we think about our work (I believe). May our shared sense of humour and joint passion for extracurricular pursuits last this lifetime.

The true (unsung) heroes of a dissertation of this nature are not the investigators but the partially unwilling subjects who serve their life's purpose aiding humankind, whether tangibly by enabling development of clinically-relevant products, or by quenching our species' thirst for knowledge. If it's not already clear, I'm referring to the non-human primate subjects used in biomedical and neuroscience research. Without them, even a small subset of this work would be inconceivable. I would like to thank the monkeys I have encountered during the course of my work, whether they ended up contributing to the data that are in here or not, because each encounter was a unique and rewarding learning experience. Waylan, Gizmo, Simon, Garfunkel, Wilma, Nova, and most recently, Blog, have all contributed immensely to my (and our) understanding of the working of the mysterious soft matter substance called the brain. But the true champion in this process has been a special monk called BB, who started his research career at the same time as me, and has grown to be a source of wonder and inspiration for me. The smartest of the lot, BB has been my go-to testbed for any new experimental ideas, and has almost always immediately provided me though-provoking data due to his ability and willingness to quickly adapt to new

paradigms. We have a special connection, as anyone who has seen us in the "wild" can attest. He is one non-human I immensely respect and adore and will be sad to bid farewell to after my stint here.

Finally, I would like to express my love for and gratitude towards Ashwati, my then fiancée and now wife (yes, that is how much I have delayed this) for her love and support during the home stretch of my dissertation pursuits. Funnily enough, I'm indebted to this dissertation itself when it comes to her, for it was during an impromptu "dissertation writing session" (more than a year before the actual writing started) that we first really *met*, curiosity meeting clarity, after four years of acquaintance and fortuitous happenstance purely remaining to serve as hippocampal fodder. We journeyed through the final laps together through music and inane laughter, long walks, deep silences, and deeper debates, through fall colours and snow, winding drives and musical surprises, balanced pragmatism meeting thoughtful radicalism, wonderfully timed gastronomic salves and endless cups of tea (for body and soul), and insidious-yet-innocuous pokes (sorry), and I stand enriched for the experience. Her passion for teaching is insatiable, and so much more remains to be learnt. This long overdue wrap-up is completely owed to her (apart from procedural requirements).

Here's to many more dissertations in life. Aum tat sat.

1.0 GENERAL INTRODUCTION – IMPORTANCE OF GAZE CONTROL

The visual environment is filled with objects that contain potentially useful information for the survival of an organism. In order to best extract this information, the animal needs to orient itself to the part of the visual world most relevant for its immediate behavior while ignoring the unwanted parts. This is especially true for foveating animals such as humans and other primates, where only a small part of the retina (the fovea) is able to resolve the visual world with greatest detail. In order to orient the fovea appropriately, we redirect our line of sight (or gaze) onto the object or spatial location of interest. The simple act of looking can provide us with information about, for example, whether the looming image in our peripheral vision is a conspecific or a predator (e.g., tiger) / dangerous object (e.g., moving vehicle), whether the expression on a person's face is anger or surprise, where to reach to acquire food, and what comes next in this sentence. However, the computations the brain needs to perform before initiating a gaze shift are not as simple: the incoming stream of sensory stimuli must be parsed in order to decide whether a gaze shift is appropriate and where the next gaze shift should be to, gaze at the current location must be disengaged, the movement must be planned, and the evolving decision must be relayed to movement generators at the appropriate time in the form of a movement command to actuate the gaze shift. Importantly, although the description above goes through a series of stages, it is also possible that they are implemented in a parallel fashion in the ongoing activity in the brain.

1.1 GAZE CONTROL NETWORK IN THE PRIMATE BRAIN

Gaze shifts (or saccades) are controlled by an extensive, distributed network of brain regions (Wurtz et al., 2001) that together comprise the oculomotor system. Incoming visual stimuli are processed by the visual cortices and visually responsive neurons in the dorsal fronto-parietal

network - that includes the lateral intraparietal area (LIP) and the frontal and supplementary eye fields (FEF and SEF) - and the midbrain superior colliculus (SC). Following selection and resolution of potential saccade targets, neurons in the FEF and SC release motor programs that activate burst generators in the brainstem reticular formation, which in turn issue the final movement command to generate a saccade. The work is this thesis primarily focuses on the SC, owing to its key role in mediating the sensorimotor transformations leading to saccades. As such, its anatomical and physiological properties are described below.

1.1.1 The Superior Colliculus

The SC is an evolutionarily ancient structure whose main function seems to be to direct or orient the attention of an animal, primarily by controlling its gaze. Located at the roof of the brainstem, the SC can be anatomically divided into seven distinct layers (for in depth reviews, see Huerta and Harting, 1984; Sparks and Hartwich-Young, 1989; May, 2006; Isa and Hall, 2009). These layers can be further classified based on their physiological and functional properties into two levels. The superficial layers (SCs) receive inputs directly from the retina as well as primary and extrastriate visual cortices (V1, V2, V4). The pretectum and parabigeminal nucleus (the cholinergic isthmic nucleus in non-mammals) also project to SCs. SCs in turn projects ventrally to the intermediate and deep layers in the SC, the lateral geniculate nucleus (LGN) in the thalamus, and reciprocally to the pretectal and parabigeminal nuclei. The intermediate and deep layers (SCid; henceforth just called intermediate layers) receive inputs from the superficial layers, the (dorsal) fronto-parietal cortical network involved in the processing of visuo-spatial information, including the lateral intraparietal area (LIP) and the frontal eye fields (FEF), the dorso-lateral prefrontal cortex (dlPFC), and the infero-temporal cortex (IT) in the ventral visual pathway. The basal ganglia also project to the SCid, primarily in the form of GABA-ergic projections from the substantia nigra pars reticulata (SNpr).

An important property of the SC network, across layers, is that each colliculus contains a topographic representation of the animal's contralateral hemifield in a retinotopic reference frame (see review by Gandhi and Katnani, 2011). Accordingly, neurons in the superficial layers exhibit responses that are time locked to the onset of visual stimuli (visual neurons) in their receptive field.

Neurons in the intermediate layers fire in response to gaze shifts (movement neurons) into their movement field, or both to a gaze shift and visual stimulus onset in their response field (visuomovement neurons). The response field locations, to a large extent, overlap across the two layers as one proceeds dorsoventrally through the SC (similar to a cortical column). The topography on the SC tissue is as follows. The rostral half of the SC maps onto the central visual field or small stimulus eccentricities and gaze shift amplitudes. The caudal half maps onto the peripheral visual field—eccentric locations and large movements. Similarly, the lateral halves of the colliculi represent the downward hemifield while the medial halves represent the upward hemifield. The mapping from visual space to SC tissue space is nonlinear, and in fact logarithmic, along the rostrocaudal axis such that significantly more neural tissue is dedicated to the central field compared to the extremities. This rich diversity in the physiological properties of the SC, its anatomical location, and its multi-modal, quasi-columnar organization supports the idea that it is a critical node in the process of integrating sensory information to produce overt orienting behavior (e.g., saccades).

1.2 SENSORIMOTOR INTEGRATION IN GAZE CONTROL

The brain is primarily a sensorimotor organ - it receives information from the environment through the senses, processes this information, and guides actions in accordance with the needs of the organism. In the gaze control system, this reduces to sorting through the rich clutter that is the visual world to decide where to look next, and when. Despite numerous advances over the last few decades, it is unclear precisely when and how this decision is made and is an important unsolved problem in systems neuroscience. However, the wealth of accumulated knowledge has contributed significantly towards identifying the potential neural substrates mediating sensorimotor transformations for gaze shifts. For example, across the gaze control network, there exist neurons that exhibit phasic activity in response to a visual stimulus (visual neurons), neurons that burst around the time of a gaze shift (movement neurons), as well as neurons that show both types of responses (visuomovement neurons) (Sparks and Mays, 1990; Wurtz et al., 2001). The latter two types of neurons, in FEF and SC, also project directly to the saccade-generating circuitry in the brainstem (Segraves, 1992; Rodgers et al., 2006). Note that the transformation from visual input

to saccadic eye movement does not happen immediately except in the rare case of express saccades (Dorris and Munoz, 1995). During the intervening period (one that is presumably important for efficient sensorimotor integration), fixation on the previous gaze location is maintained. Control of fixation and fixational eye movements is enabled by a dedicated subnetwork distributed across FEF, SC and the brainstem (Cohen and Henn, 1972; Munoz and Wurtz, 1993a; Hafed et al., 2009; Izawa et al., 2009). Reciprocal interactions between fixation-related neurons and the saccade-generating circuitry have been reported (Munoz et al., 1996; Munoz and Istvan, 1998; Gandhi and Keller, 1999); however, the role of fixation in enabling sensorimotor transformations has not been thoroughly studied.

In addition, many visuomovement neurons exhibit a low level of activity in the period between sensory and movement-related transients. This low-frequency discharge and modulations thereof have been attributed to many functionally distinct processes, including but not limited to motor preparation (Bracewell et al., 1996; Hanes and Schall, 1996; Mazzoni et al., 1996; Dorris et al., 1997; Dorris and Munoz, 1998), attention (Goldberg and Wurtz, 1972; Colby and Goldberg, 1999; Ignashchenkova et al., 2004; Thompson et al., 2005; Buschman and Miller, 2007, 2009), working memory (Balan and Ferrera, 2003; Zhang and Barash, 2004; Curtis and D'Esposito, 2006), target selection (Schall and Hanes, 1993; Basso and Wurtz, 1997; Horwitz and Newsome, 1999; McPeek and Keller, 2002; Sato and Schall, 2003; Carello and Krauzlis, 2004), action selection or decision-making (Newsome et al., 1989; Gold and Shadlen, 2007), target expectation (Dorris and Munoz, 1998), and reward expectation (Platt and Glimcher, 1999; Hikosaka et al., 2006; Hikosaka et al., 2008). The multiplicity of functions these neurons seemingly take part in raises the intriguing question - at what point does the information embedded in their activity go from representing a perceptual and/or cognitive process to a movement-related signal?

1.2.1 Fixation and Saccade Generation

Early studies of the neural basis of fixation reported the presence of neurons in the rostral portion of SC that seemed to mediate gaze withholding or the maintenance of fixation (Munoz and Wurtz, 1993a). These so-called fixation neurons fire at a tonic rate when the eyes are fixated on a spot and lower their activity during saccades, when neurons in caudal SC show elevated firing. Microstimulation of the rostral SC stops saccades in mid-flight (Munoz et al., 1996), and can also

suppress buildup activity in the caudal region (Munoz and Istvan, 1998). These observations gave rise to the fixation zone–saccade zone model, where long-range reciprocal inhibition between the rostral and caudal "zones" control the alternating pattern of fixations and saccades seen during typical gaze behavior.

However, several lines of evidence have since emerged that dispute this model. First, the nature of perturbations in saccades caused by rostral SC microstimulation fall along a continuum, along with those caused by stimulation of caudal SC, and are qualitatively different from interruptions caused by stimulation of the omnipause neurons (OPNs) in the PPRF (Gandhi and Keller, 1999). The OPNs exhibit activity that is much more closely linked to the onset and offset of fixation and therefore make a more qualified candidate for withholding gaze. Second, experiments in the cat employing multi-step gaze shifts have shown that the locus of activity across SCid is correlated with the distance to the eventual goal of the sequence of gaze shifts - or goal-based long term motor error - with rostral SCid neurons returning to their tonic firing only after the final target is fixated (and not during intermediate fixations) (Bergeron et al., 2003). Third, recent work has demonstrated a causal role for rostral SC neurons in the generation of microsaccades - tiny movements (<1') of the eyes during fixation (Hafed et al., 2009). These findings are in line with the idea that the SC map represents a natural continuum of movements, with large movements represented toward the caudal end and small movements represented near the rostral end. Thus, the transition from withholding gaze to shifting gaze and vice versa is better explained by a model that considers the shifts of balance between the rostral portion of SC encoding small movements (or target errors) and the caudal portion encoding large movements. The maintenance of this balance presumably prevents undesirable activity leading to premature movements under normal conditions (Meredith and Ramoa, 1998) and enables sensorimotor transformations leading to saccades.

1.2.2 Movement Preparation and Initiation

As mentioned above, many neurons in SC exhibit low frequency activity between the onset of a stimulus and a gaze shift to that stimulus, and studies with cleverly designed behavioural tasks have assigned this activity to diverse range of perceptual and cognitive functions. Of course, the typical SC neuron discharges a high-frequency burst prior to a saccade into its movement field,

and it also projects to the brainstem burst generator elements that execute saccades. Thus, studies have proposed a preparatory premotor component to the low-level response (Glimcher and Sparks, 1992; Hanes and Schall, 1996; Mazzoni et al., 1996; Dorris et al., 1997; Steinmetz and Moore, 2010) and results assigning a cognitive role to the neural discharge have been subject to the criticism that the low-level activity may instead reflect a preparatory command for a movement (usually a saccade) that is planned but not necessarily executed (Gandhi and Sparks, 2004; Ignashchenkova et al., 2004; Krauzlis et al., 2004). Of course, it is also likely that both cognitive and premotor signals are reflected in the low-level discharge, a hypothesis that is the basis of the "premotor theory of attention" (Rizzolatti et al., 1987). While links between visual attention and saccade preparation have been implied by psychophysical (Rizzolatti et al., 1987; Hoffman and Subramaniam, 1995) and neurophysiological (Corbetta et al., 1998; Moore and Fallah, 2001; McPeek and Keller, 2002; Cavanaugh and Wurtz, 2004; Muller et al., 2005; Awh et al., 2006) studies, evidence also exists for distinct attention and preparation processes, particularly at the level of FEFs (Juan et al., 2004; Schall, 2004; Thompson et al., 2005; Gregoriou et al., 2012).

The motor preparation hypothesis states that the low-level discharge in gaze control network neurons accumulates gradually toward a cell-specific threshold, at which point (a) it converts into a high-frequency burst, (b) the brainstem OPNs become quiescent, and (c) a saccade is triggered (Hanes and Schall, 1996; Dorris et al., 1997). It has been hypothesized that the low frequency discharge represents a movement preparation signal that encodes both timing and metrics of the desired saccade. Indeed, the firing rate level in the preparatory period is negatively correlated with the saccade reaction time (the higher the activity, the earlier the movement occurs), and the locus of activity in the SC dictates the saccade vector. Basic computational models that simulate the accumulation or drift rate as either noisy (Ratcliff et al., 2003; Lo and Wang, 2006) or ballistic (Carpenter and Williams, 1995; Reddi et al., 2003) can sufficiently describe the trial-to-trial variability in the discharge patterns of SC neurons and the distribution of reaction times. This functional assessment of movement preparation, however, is gauged by correlating neural activity with movement features that are observed hundreds of milliseconds later, after the animal is granted permission to generate a response. A stronger foundation for movement preparation, and its time course, could be established if a behavioral output can be revealed as the low-frequency activity is evolving.

It is important to mention here that motor systems research outside the field of oculomotor control has followed a different trajectory to understand the dynamics of movement preparation and initiation. Studies of movement preparation in the skeletomotor system have shown that neural activity reaches an optimal subspace before undergoing dynamics that produce a limb movement (Afshar et al., 2011; Churchland et al., 2012). A related idea suggests that activity pertaining to movement preparation evolves in a region of population space that is orthogonal to the optimal subspace, and this dissociation confers neurons the ability to prepare the movement by incorporating perceptual and cognitive information without risking a premature movement (Kaufman et al., 2014).

1.3 RESEARCH OBJECTIVES

While the field of oculomotor research is one of the oldest in neuroscience and the community has made significant strides in understanding the neural bases of gaze control and associated sensory-to-motor transformations, several fundamental questions remain unanswered to a unanimous level of satisfaction. The following sections summarize the several unknowns we have mentioned in this Introduction thus far, and describe our approach(es) to resolve them.

1.3.1 Knowledge Gap

There are three salient unknowns that emerge from the discussion above:

1) What role do fixation and fixational eye movements play in dictating the time course of sensorimotor transformations?
2) How do the dynamics of low frequency preparatory activity inform movement initiation? Does preparatory activity possess ‘motor potential”?
3) How do downstream decoders decide to not trigger a movement in response to the high frequency volley of inputs following onset of a visual stimulus?

1.3.2 Research Strategy

Prior to triggering the movement, the saccade generation circuitry must overcome the potent inhibition imposed to preserve fixation. A potential method to reveal the presence and evolution of movement preparation exploits the antagonistic relationship between motor preparation and saccade inhibition. Located in the pons, the OPNs inhibit the saccade burst generator circuit that innervates the extraocular motoneurons. The OPNs discharge at a tonic rate during fixation and cease activity during all saccades. Consider the scenario in which the eyes are stable so that the OPNs are active at a tonic rate, and an experimental manipulation is available to transiently inhibit them at different times after object(s) are presented in the visual periphery. This early and transient withdrawal of inhibition could, in principle, "trick" the saccadic system into prematurely executing an eye movement if the underlying low-frequency discharge in SC and the brainstem burst generator reflects a premotor signal; note that this reasoning remains agnostic about cognitive signals present simultaneously. If successful, the timing of the earliest saccade will indicate when movement preparation commences, and the kinematics will reveal how it develops. Moreover, concurrently recording activity in the SC should provide a neural correlate of timing, kinematics, and direction of prematurely triggered saccades.

Related previous work in our lab has utilized the observation that OPNs also become quiescent during blinks (Schultz et al., 2010), and delivered an air-puff to one eye to invoke the trigeminal blink reflex in non-human primates performing various saccade tasks (Gandhi and Bonadonna, 2005). A blink evoked shortly after target onset but before the typical saccade reaction time triggered a saccade of shorter latency. There was a strong correlation between blink time (removal of OPN inhibition) and saccade time, providing a behavioural readout of movement preparation. It has since also been used to probe the time course of movement cancellation (Walton and Gandhi, 2006) and motor preparation in paradigms geared to understand target selection (Katnani and Gandhi, 2013). In the latter study, they evoked blinks in monkeys performing a visual search task and demonstrated that the movement preparation signal evolves as early as the neural modulation associated with target selection.

We use reflex blinks to answer two of the three questions outlined in the previous section. For the third question, we turn to population dynamics for analytical insights

2.0 DISRUPTION OF FIXATION REVEALS LATENT SENSORIMOTOR PROCESSES IN THE SUPERIOR COLLICULUS

2.1 INTRODUCTION

Active vision entails interacting with the environment through alternating patterns of fixations and redirections of gaze or saccades. The behavioral transition between fixation and saccade generation is mediated by a neuronal shift between movement-suppressing neurons and movement-generating neurons in the oculomotor neuraxis. The superior colliculus (SC), which is a major node in the gaze control pathway (Gandhi and Katnani, 2011), is ideally situated to implement the transition. Movement suppression is controlled by so-called "fixation-related neurons" in rostral SC (Munoz and Wurtz, 1993a). These neurons are active during fixation, which is likely maintained by preparatory activity associated with programming of microsaccades (Hafed et al., 2009). Crucially, they suppress their discharge before large saccades. Movement generation is mediated by visuomovement neurons in caudal SC (Sparks and Mays, 1990; Dorris et al., 1997) that burst during saccades. Intriguingly, these neurons also burst in response to a visual stimulus.

One line of evidence points to a reciprocal relationship between rostral and caudal SC in maintaining the balance between fixation and saccade generation under normal conditions. Inactivation of rostral SC leads to uninhibited saccade generation to peripheral visual targets (Munoz and Wurtz, 1993b) or stable offsets during fixation (Hafed et al., 2008). Conversely, microstimulation in caudal SC generates saccades (Robinson, 1972), and microstimulation of rostral SC interrupts ongoing large movements (Munoz et al., 1996), perhaps through interference induced by stimulation-evoked small amplitude saccades (Gandhi and Keller, 1999). However, the inability to record neural activity simultaneously during these manipulations makes it difficult to gauge the functional relationship between fixation and saccade-generating mechanisms, and

whether the interactions extend beyond movement generation to other features of the sensorimotor transformation. For instance, does the state of the network during fixation influence the visual response of visuomovement neurons in caudal SC? Are activation changes in one network obligatorily linked to changes in the other, and how do these changes influence fixation, microsaccade and saccade behavior? It is important to mention here that we use the term "network" in this chapter to broadly refer to a collection of putatively interconnected neurons that are thought to perform a specified function.

We used an alternative approach – naturalistic, non-invasive and transient behavioral perturbation of fixation – that permitted simultaneous recording of neural activity. We reasoned that if fixation is behaviorally disrupted, it must affect activity in the fixation-related network, altering the balance between the fixation and saccade-generating networks. Hence, on some trials within a session, we presented an air puff to the subject's eye to evoke the trigeminal blink reflex (Berardelli et al., 1985) during initial fixation of a central target. Blinks disrupt fixation by transiently removing foveal visual input and producing an associated blink-related eye movement (Rottach et al., 1998; Gandhi and Bonadonna, 2005), during which activity in the pontine omnipause neurons (OPNs), a low-level node in the fixation network, is suppressed (Schultz et al., 2010). We asked whether and how this perturbation affects the evolution of activity in the fixation-related and saccade-generating networks in the SC. Note that the blink was induced while the animal directed its gaze to a central target and, crucially, fixation of this target was re-established before a peripheral stimulus was illuminated. Thus, we were able to circumvent potential confounds of transient changes caused by the blink perturbation. In a small percentage of trials, blinks biased the network-level balance towards caudal SC neurons and led to undesirable behavior in the form of premature saccades. Most importantly, we observed changes in network activity even when saccades were <u>not</u> triggered prematurely. The stimulus-evoked visual response of visuomovement neurons was enhanced, and surprisingly, putative movement neurons also exhibited a visual response. Blink-induced modulation of rostral SC activity was better associated with the animal's propensity to make microsaccades. It did not obey the purported coupled relationship with saccade-generating neurons in the caudal SC. These results have a potentially deep impact on our understanding of mechanisms mediating voluntary control of gaze and its dysfunction in neuropsychiatric disorders.

2.2 METHODS AND DATA ANALYSIS

2.2.1 General and surgical procedures

All experimental and surgical procedures were approved by the Institutional Animal Care and Use Committee at the University of Pittsburgh and were in compliance with the US Public Health Service policy on the humane care and use of laboratory animals. We used three adult rhesus monkeys (*Macaca mulatta*, 2 male, ages 8 and 6, and 1 female, age 10) for our experiments. One animal expired during the course of the study, so we completed the remaining experiments with the third animal. Under isoflurane anesthesia, a craniotomy that allowed access to the SC was performed and a recording chamber was secured to the skull over the craniotomy. In addition, posts for head restraint and scleral search coils to track gaze shifts were implanted. Post-recovery, the animal was trained to perform standard eye movement tasks for a liquid reward.

2.2.2 Visual stimuli and behavior

Visual stimuli were displayed either by back-projection onto a hemispherical dome (monkeys BB and WM), or on a LED-backlit flat screen television (monkeys BB and BL). Stimuli were white squares on a dark grey background, 4x4 pixels in size and subtended approximately 0.5° of visual angle. Eye position was recorded using the scleral search coil technique, sampled at 1 kHz. Stimulus presentation and the animal's behavior were under real-time control with a LabVIEW-based controller interface (Bryant and Gandhi, 2005). After initial training and acclimatization, the monkeys were trained to perform a delayed saccade task. The subject was required to initiate the trial by acquiring fixation on a central fixation target. Next, a target appeared in the periphery but the fixation point remained illuminated for a variable 500-1200 ms, and the animal was required to delay saccade onset until the fixation point was extinguished (GO cue). Trials in which fixation was broken before peripheral target onset were removed from further analyses. The animals performed the task correctly on >95% of the trials.

2.2.3 Perturbations

On a small percentage of trials (~15-20%), fixation was perturbed by delivering an air puff to the animal's eye to invoke the trigeminal blink reflex. Compressed air was fed through a pressure valve and air flow was monitored with a flow meter. To record blinks, we taped a small Teflon-coated stainless steel coil (similar to the ones used for eye tracking, but smaller in coil diameter) to the top of the eyelid. The air pressure was titrated during each session to evoke a single blink (typically 15-30 psi). Trials in which the animal blinked excessively or did not blink were aborted and/or excluded from further analyses. Air puff delivery was timed to evoke blinks 400−100 ms before target onset. Blinks typically lasted 50−100 ms, and we removed trials in which fixation of the central target was not re-established before a stimulus was illuminated in the periphery to ensure that vision of fixation spot or the incoming stimulus was not altered at the time of its onset. For two of the monkeys (BB and BL), we also used two other forms of perturbation. First, the central fixation target was transiently blanked out for 50−100 ms at approximately the same times that a blink would have occurred on blink trials (before peripheral target appearance). We refer to these as target blank trials. In other experiments, an air puff was introduced to the animal's ear. Trials in which the animal blinked defensively in response to the ear puff were removed from further analysis. The animals typically adapted to this perturbation after a few trials.

2.2.4 Electrophysiology

During each recording session, a tungsten microelectrode was lowered into the SC chamber using a hydraulic microdrive. Neural activity was amplified and band-pass filtered between 200 Hz and 5 kHz and fed to a digital oscilloscope for visualization and spike discrimination. A window discriminator was used to threshold and trigger spikes online, and the corresponding spike times were recorded. The location of the electrode in the SC was confirmed by the presence of visual and movement-related activity as well as the ability to evoke fixed vector saccadic eye movements at low stimulation currents (20-40 μA, 400 Hz, 100 ms). Before beginning data collection for a given neuron, its response field was roughly estimated. During data collection, the saccade target was placed either in the neuron's response field or at the diametrically opposite location (reflected across both axes) in a randomly interleaved manner. For recordings in rostral SC, stimuli were

presented at one of two or four locations at an eccentricity sufficient to induce a reduction in activity during the large amplitude saccade. In two of the monkeys (BB and BL), we also performed recordings with a 16-channel linear microelectrode array (Alpha Omega, Inc., 150 micron spacing between electrode contacts). Neural activity was digitized, filtered, and recorded with the Grapevine Scout system (Ripple, Inc.). Since the electrode's approach was normal to the SC, neurons recorded with the array typically had comparable response fields.

2.2.5 Data analysis and pre-processing

Data were analyzed using a combination of in-house software and Matlab. Eye position signals were smoothed with a phase-neutral filter and differentiated to obtain velocity traces. Saccades were detected using standard velocity criteria. The animal was considered to be maintaining fixation if the gaze remained within a 2-3° window around the fixation target. Slow, blink-related eye movements were detected using lower velocity criteria. We also detected microsaccades that occurred during a trial by using a velocity criterion based on the overall variability in the velocity signal for that trial – the adaptive algorithm allowed us to account for changes in fixation patterns and movements over time for each subject. We visually verified that microsaccades were being detected accurately.

Raw spike density waveforms were computed for each neuron and each trial by convolving the spike trains with a Gaussian kernel (width = 4 ms for caudal SC, 10 ms for rostral SC). For a given neuron and target location, spike densities were averaged across trials after event-related alignment (perturbation/target/saccade onset). Caudal SC neurons were classified as task-related (visual and/or movement) if the activity during the visual epoch (100-250 ms following target onset) and/or the premotor epoch (-100 to 50 ms around saccade onset) was significantly elevated above baseline. Rostral SC neurons were classified as fixation-related if their activity during the premotor epoch of large saccades was significantly reduced below baseline. We used a firing rate criterion to increase confidence that we were using only meaningful neurons in our analysis. This was especially important for the multi-channel recordings as almost all channels had some activity due to crossover spikes or noise. Hence, we only considered caudal SC neurons that had a maximum firing rate of at least 100 spikes/s and rostral SC neurons that had a tonic baseline firing

rate of at least 20 spikes/s. Further, we only used neurons which had at least 5 trials for a given condition (e.g., early saccades following blink perturbation). Since the animals rarely made errors in the control condition, we did not have a sufficient number of neurons that met this criterion for that condition to warrant further analysis. To enable comparison of spike densities during the perturbation epoch, for no-perturbation trials, the perturbation-related alignment was performed by random assignment of "perturbation" times from the distribution of times from perturbation trials for that session. We also normalized the trial-averaged spike density of each neuron to enable meaningful averaging across the population. The activity of each caudal SC neuron was normalized by its peak firing rate across conditions. The activity of each rostral SC neuron was normalized by the baseline firing rate across conditions.

Statistical comparisons of neural activity between conditions were performed using non-parametric tests (Wilcoxon rank-sum and signed-rank tests). Activity comparisons wither used the average activity in a window, or the instantaneous activity (sliding 1 ms bin), as appropriate. Where applicable, multiple comparisons were corrected for (Bonferroni correction). A p-value < 0.01 was used to assess statistical significance, unless otherwise specified.

To quantify the functional properties of caudal SC neurons as visual, visuomovement, or movement-related, we computed the visuomovement index (VMI) for each neuron. VMI was calculated as $VMI = \frac{V-M}{V+M}$, where V is the baseline-subtracted activity during the visual epoch, and M is the delay period-subtracted activity during the saccade epoch. We also calculated a modulation index to quantify the perturbation-induced modulation in the neurons' visuomovement activity. The modulation index was calculated for the visual epoch as $\frac{V_p-V_c}{V_p+V_c}$, where V_p and V_c are the baseline-subtracted activity during the visual epoch in perturbation and control trials, respectively.

2.3 RESULTS

We perturbed fixation by inducing a blink and recorded neural activity from the SC in three monkeys performing a delayed response task (Figure 21a). During the control condition, the animals correctly initiated the eye movement after fixation point offset on 98.0% percent of the trials. The black trace (see "Eye" label, Figure 21a schematic) illustrates the temporal profile of a typical saccade and the black histogram (Figure 21b) shows the distribution of reaction times (mean +/- s.d = 276 +/- 46 ms). On perturbation trials, the blink was always induced during the initial fixation period (small bumps in the colored eye position traces above the red arrows, Figure 21a). We only analyzed trials in which the blink had completed and fixation re-established before a target was illuminated in the visual periphery.

For the majority of trials (4372 out of 5178; 84.4%), the animals produced a saccade after the saccade initiation cue, much like the behavior seen during control trials. The red trace (Figure 21a) represents a typical example and the same color histogram (Figure 21b) shows the associated latency distribution (mean +/- s.d = 271 +/- 50 ms). On a smaller subset of trials (806 of 5178; 15.6%), a saccade was produced prematurely during the delay period. The gold trace (Figure 21a) denotes a typical example and the same color histogram (Figure 21b) displays the latency distribution. The increase in likelihood of premature saccades was observed not just across the entire dataset but also during a majority of sessions in every animal (Figure 21c; Wilcoxon signed-rank test, $n=89$, $p<0.0001$). Neural activity recorded during blink trials that did not alter behavioral performance ("red" data) provide an opportunity to assess the latent functional interactions between fixation and saccade-generating networks without introducing any confounds from activity associated with saccade execution. Analyses of these data constitute the major focus of this study. Discharge patterns associated with the smaller subset of blink data that yielded premature saccades ("gold" data) are considered briefly at the end of the Results section.

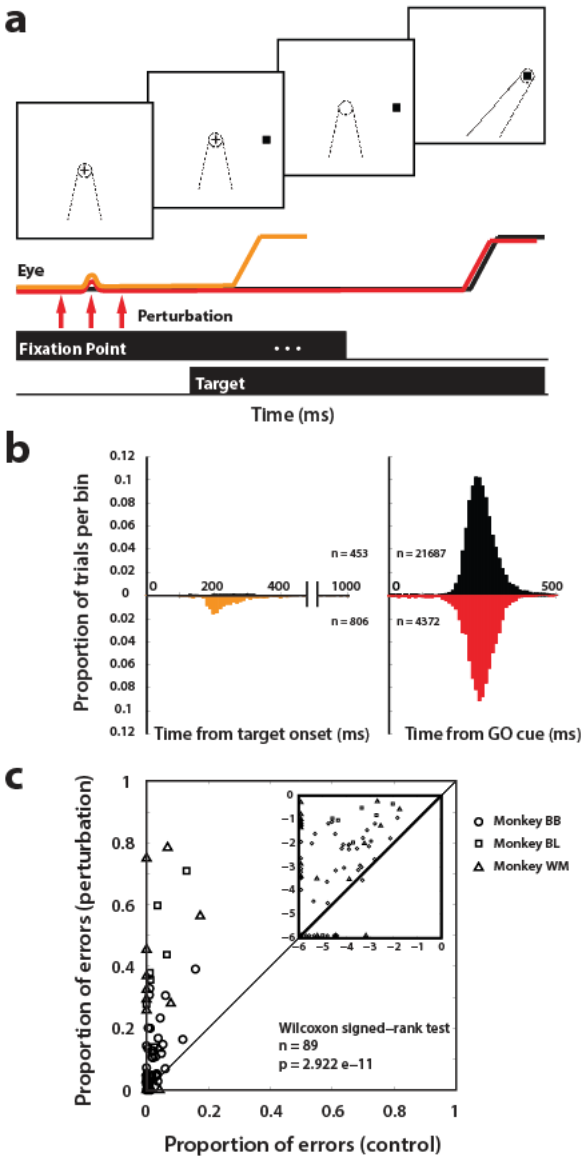

Figure 21. Behavior in control and perturbation conditions

a. Schematic of the delayed saccade task. The red arrows indicate that blink perturbation times preceded target appearance. The temporal traces in the bottom rows are schematics of typical eye position profiles in the absence of perturbation (control, black) and with perturbation leading to premature (gold) and regular latency saccades (red). Radial position is plotted, hence all deflections are positive. **b.** Histograms for the distribution of saccade reaction times with regular (right) and early (left) latencies for control (top) and perturbation (bottom) trials. The offset of fixation point represents the "GO" cue. **c.** Proportion of errors in the perturbation condition (ordinate) plotted against the control condition (abscissa). Each point represents a session. Symbols represent individual subjects. The unity line is on the diagonal. The inset replots the same data on log-log axes for better visualization.

2.3.1 Sub-threshold enhancement of stimulus-evoked activity following disruption of fixation

Figure 22a shows the population activity of 94 caudal SC neurons recorded in the intermediate layers during control (black) and perturbation (red) trials with regular latency saccades. The visual response was strongly enhanced at the population level (Wilcoxon signed-rank test, $p<0.01$), as indicated by the significance bands along the bottom of the figure. Note that the perturbation occurred before target onset and the eyes were stationary and re-fixated on the fixation point at the time of the visual burst (bottom row in Figure 22a shows mean+/-s.e.m eye position traces for each condition). Activity was also enhanced significantly into the delay period. Scatter plots in Figure 22b illustrate that the increase was observed irrespective of the neurons' firing rate, and in several cases, the difference between conditions was of the order of 100 spikes/s or more. This suggests that the perturbation-induced changes to the saccade-generating circuitry were strong and long-lasting, persisting several hundreds of milliseconds into the trial (left and middle panels). In contrast, the activity during the motor burst was comparable between control and perturbation conditions (right panel in Figure 22b).

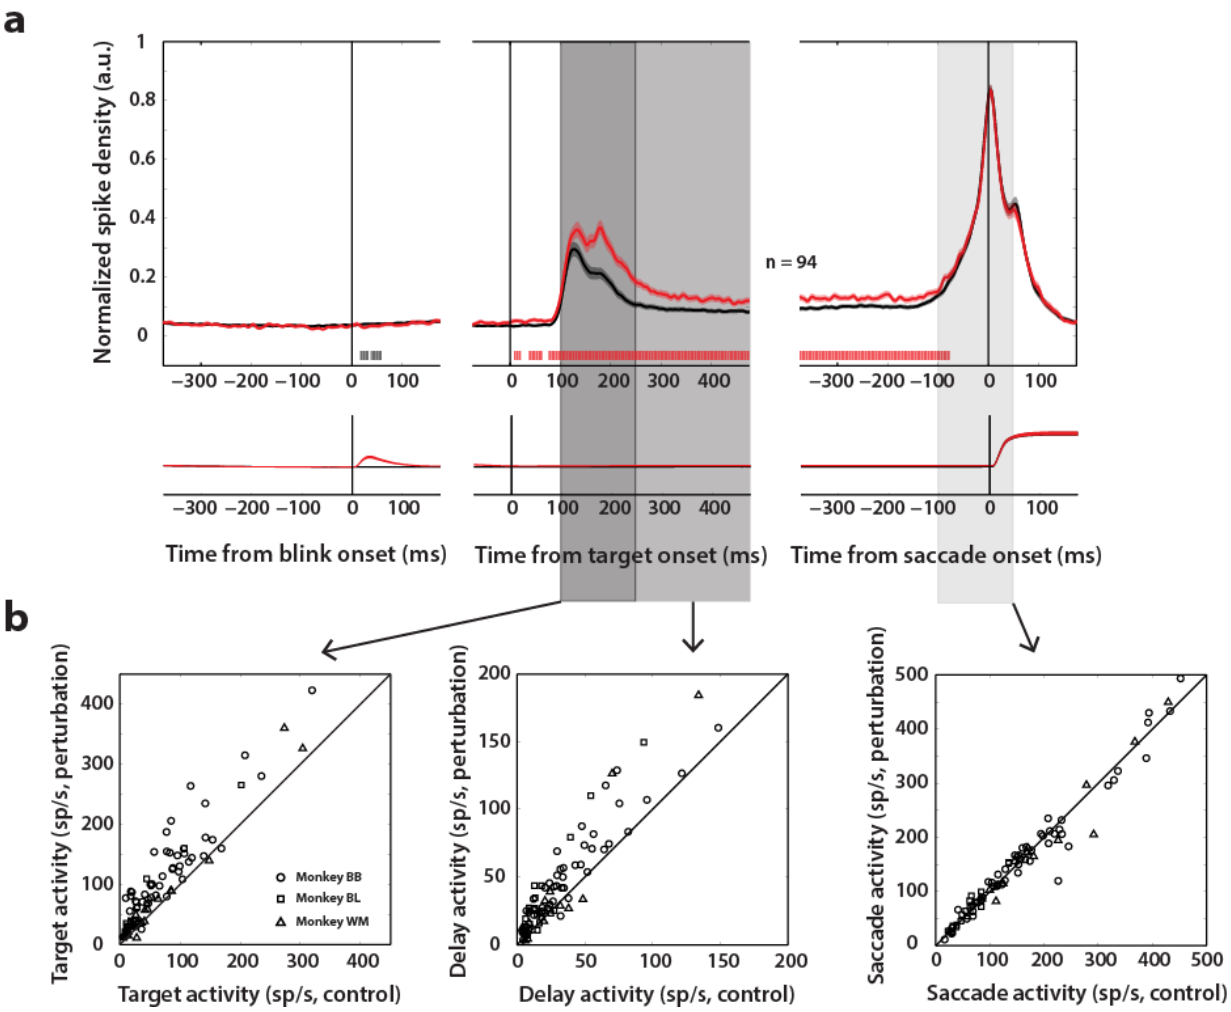

Figure 22. Population activity in caudal SC

a. Population average (mean +/- s.e.m.) of caudal SC neurons (normalized) for the control (black) and perturbation (red) conditions. The tick marks near the base of the figure indicate time points at which the difference between the two conditions was significantly different (Wilcoxon signed-rank test, $p<0.01$). The color of the tick mark indicates the condition in which activity was higher. The bottom row shows mean eye position traces in the two conditions. **b.** Scatter plots of individual neurons' activities (n = 94) during the target-evoked response period (left, dark shaded epoch in **a**), delay period (middle, intermediate shaded epoch in **a**), and saccade period (right, light shaded epoch in **a**) in the perturbation condition plotted against the control condition. Unity line is on the diagonal. Each symbol corresponds to a different monkey. Activity was enhanced during the target and delay epochs but not during the saccade epoch for a majority of neurons.

2.3.2 A hidden visual response is present in putative movement neurons

Next, we studied how the disruption of fixation affected the functional properties of these neurons. The visuomovement index (VMI; see Methods), a measure of the relative strengths of a neuron's visual and saccade-related bursts, has previously been used to characterize SC and FEF neurons into functional classes (e.g., Shen and Pare, 2007; Cohen et al., 2009). Negative values (closer to -1) indicate that the neuron is primarily visually driven and positive values (close to +1) indicate that the neuron is primarily movement-related, with values in between representing a combination of the two processes. We computed the VMI for each neuron based on its activity on control trials and perturbation trials separately. Figure 23a shows the two sets of VMIs plotted against each other. The individual points largely fall below the unity line (Wilcoxon signed-rank test, n=103, $p<0.0001$), indicating a shift towards a visually-driven profile in the neurons' response (see also histogram). Note that several canonical movement neurons based on VMI (right most grid column) transform into visuomovement neurons in the perturbation condition, unmasking the presence of latent "visual" activity throughout the intermediate/deeper layers. The top and bottom panels of Figure 23b respectively show the population average of 11 movement neurons (VMI > 0.9) as well as the activity of an example movement neuron in the two conditions.

We then looked at whether a neuron's functional classification, as defined by the VMI, has an influence on its susceptibility to the perturbation. We computed a modulation index for each neuron based on the relative change in its activity during the visual epoch on perturbation trials compared to control trials (see Methods). The modulation index is plotted against the VMI for each neuron in Figure 23c. The two indices were highly correlated (Pearson's correlation coefficient = 0.571, $p<0.0001$), indicating that neurons with a more movement-related profile (VMI closer to +1) are modulated to a greater extent by disruption of fixation. These results suggest that the contribution of individual neurons to visuomotor transformations and saccade generation is not fixed but flexibly modulated depending on the state of the network.

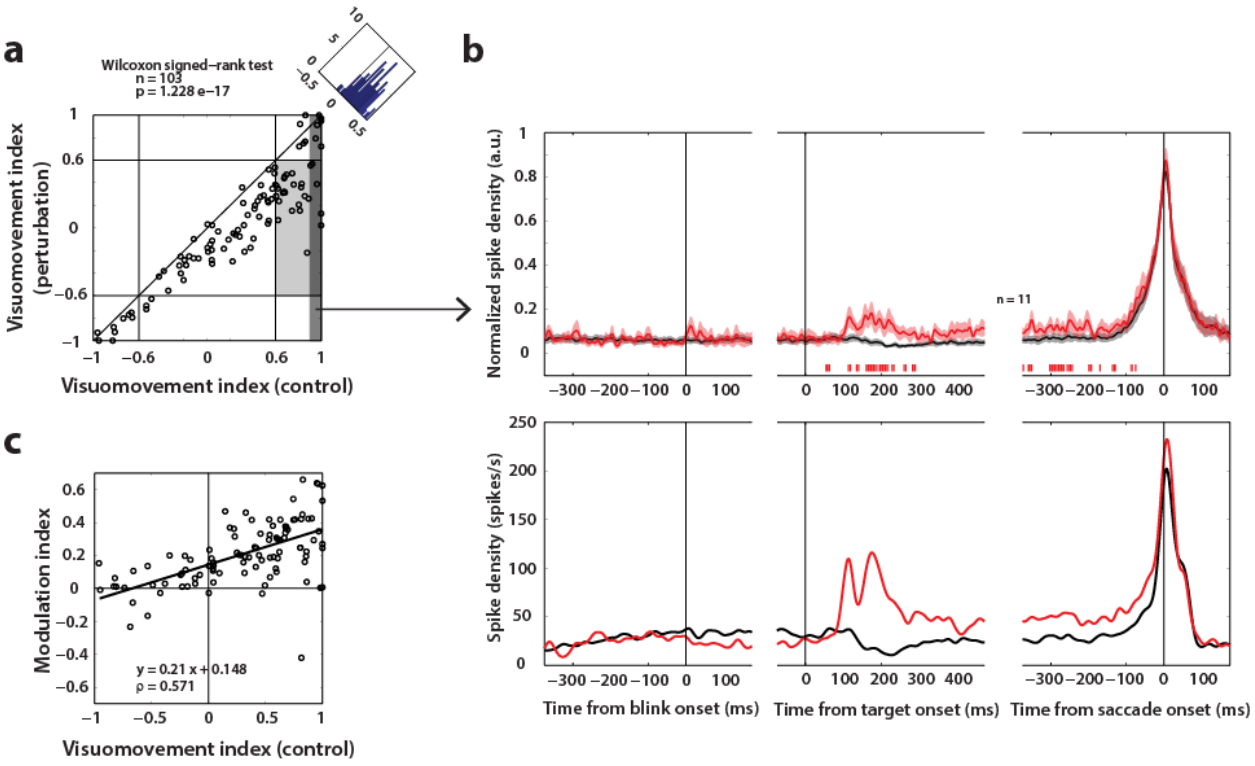

Figure 23. Latent visual response in putative movement neurons

a. Visuomovement index (VMI, see Methods) calculated from perturbation trials plotted against VMI from control trials. VMI significantly decreased in perturbation trials (neurons became more "visual"). Note that neurons in the light shaded rectangle that would typically be classified as "movement" neurons based on activity in the control condition (VMI > 0.6) become visuomovement neurons under perturbation. The histogram shows the distribution of VMI differences between the two conditions. **b.** Top. Population activity of 11 movement neurons, defined as those with VMI > 0.9 (dark shaded region in **a**), in control and perturbation conditions. Bottom. Activity of an exemplar movement neuron showing strong unmasking of a latent visual response following the perturbation. **c.** Modulation index in perturbation trials relative to control trials plotted as a function of VMI in control trials. The correlation was significantly positive, i.e., more movement-like cells were more likely to be modulated by the perturbation during the visual epoch.

2.3.3 Fixation-related activity in rostral SC is dissociated from visuomovement activity

Rostral SC neurons, which fire at a tonic rate during fixation and burst for microsaccades, are suppressed during larger saccades, and there is some evidence for reciprocal inhibitory connectivity between rostral and caudal SC (Meredith and Ramoa, 1998; Munoz and Istvan, 1998; but see also Phongphanphanee et al., 2014). We therefore considered whether the enhancement of visuomovement activity following perturbation in caudal SC is obligatorily linked to a suppression of activity in the fixation-related network in rostral SC. Figure 24a shows the population activity of 44 rostral SC neurons in the control and perturbation conditions. At the population level across

all three animals, there was no consistent difference between the two conditions as seen from the significance bands below the activity traces (Wilcoxon signed-rank test, p>0.05). However, the individual subjects' data presented a spectrum of effects (Figure 24b). Monkey BB's rostral SC showed a dramatic and sustained suppression on perturbation trials whereas monkeys BL and WM showed a transient increase immediately after the perturbation before a return back to control levels. The heterogeneity of the effect can also be realized through the scatter plot of average firing rates during the delay period (Figure 24c). Since the behavior and caudal SC profiles are largely consistent across subjects, while the rostral SC activity is not, it follows that the rostral SC cannot be exclusively responsible for the shift in balance from fixation to saccade generation.

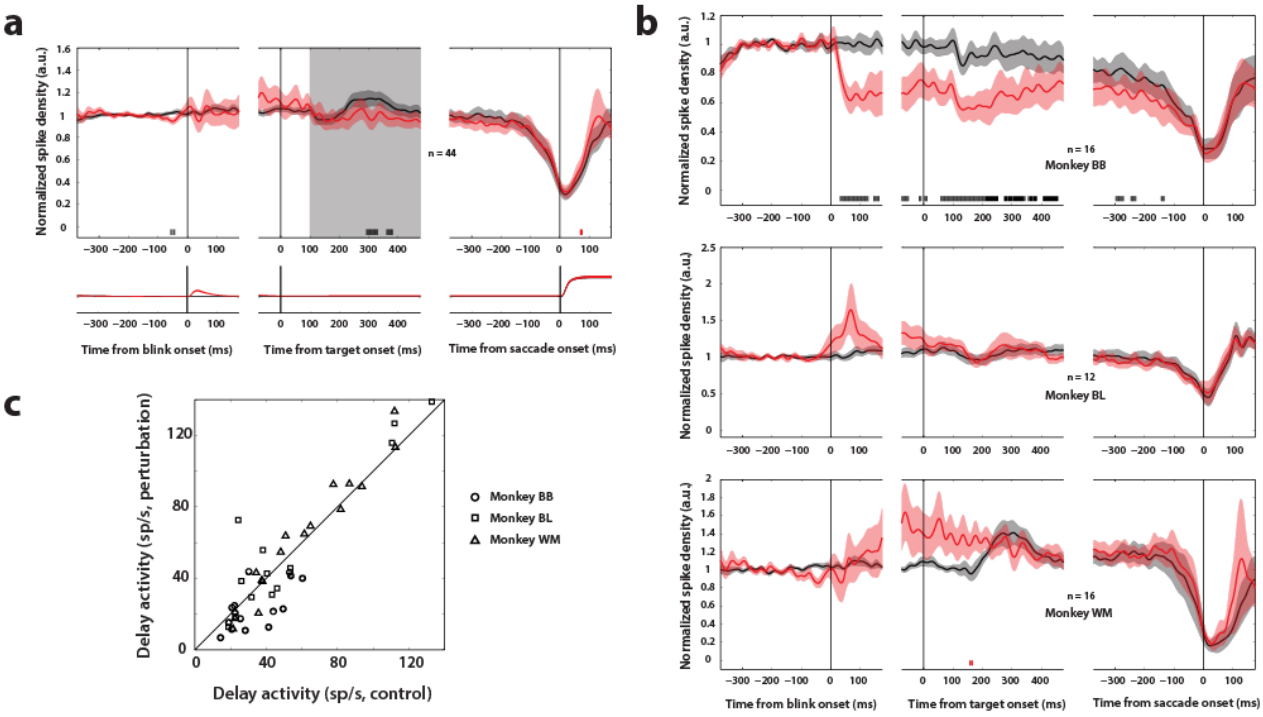

Figure 24. Population activity in rostral SC

a. Top - Population average of rostral SC neurons (normalized) in the two conditions. Colors and tick marks as in Figure 22. The bottom row shows mean eye position traces. **b.** Individual subject population means of rostral SC neurons. Activity was strongly suppressed in monkey BB whereas monkeys BL and WM showed transient increase in rostral SC activity (not significant, Wilcoxon rank-sum test, p>0.05). **c.** Scatter plot of the individual neurons' activities (n = 44) during the delay period (shaded epoch in **a**) in the perturbation condition plotted against the control condition. Plot follows the scheme in Figure 22b.

2.3.4 Rostral SC activity is better correlated with microsaccade occurrence

We asked whether another aspect of the animals' gaze-related behavior may help explain the heterogeneity in rostral SC and the apparent disconnect between the caudal and rostral SC populations. Since rostral SC neurons are known to play a role in the generation of tiny fixational eye movements known as microsaccades (Hafed et al., 2009), we looked at the microsaccade rate for the individual subjects as a function of time on control and perturbation trials (Figure 25a). Monkeys BL and WM (middle and right columns) exhibited a characteristic microsaccade signature on control trials, including target-related suppression of microsaccades followed by a rebound (Winterson and Collewijn, 1976; Engbert and Kliegl, 2003). On perturbation trials, there was a pre-target reduction in microsaccade rate for these two subjects, which we attribute to the inability to detect microsaccades during the blink-related eye movement. On the contrary, we propose that the actual microsaccade rate during this period increases, a claim consistent with the increase in rostral SC activity, particularly for monkey WM (Figure 25b). We suggest that a blink-related reduction in brainstem omnipause neuron activity removes inhibition on the burst generators (Gandhi and Bonadonna, 2005; Schultz et al., 2010) and allows activity in upstream premotor networks (in this case, fixation activity in rostral SC) to generate the associated movement (microsaccades) in these two subjects. On the other hand, monkey BB, who showed a strong and sustained suppression in rostral SC activity in the perturbation condition, rarely ever made microsaccades in either condition. Thus, microsaccades explain some of the discrepancies observed in rostral SC activity.

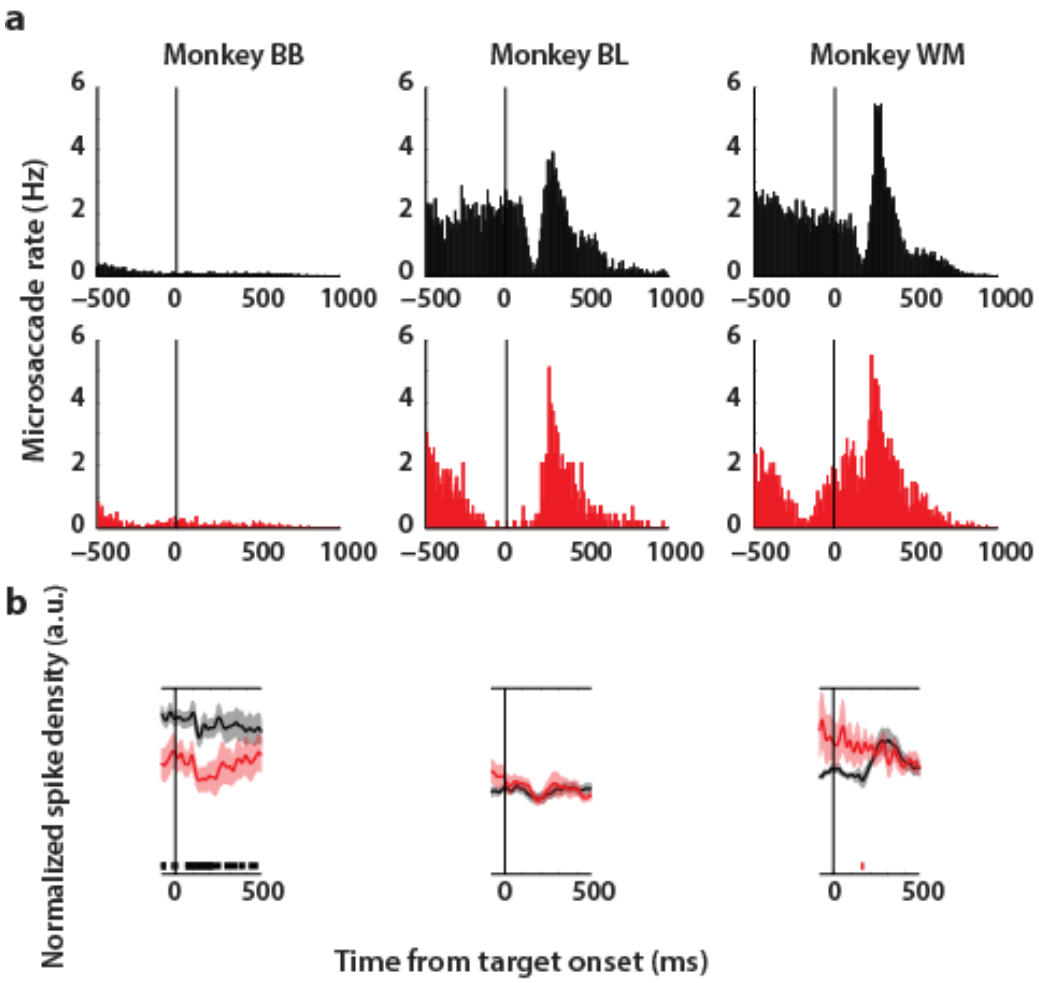

Figure 25. Microsaccade behavior of individual animals

a. Microsaccade rate as a function of time relative to target onset for each of the three monkeys in control (top row) and perturbation (middle row) trials. Monkey 1 (BB) rarely made microsaccades at any time after acquiring fixation of the central fixation target. The other two monkeys show a characteristic microsaccade rate profile with transient inhibition of microsaccades following target onset. The dramatic reduction in microsaccade occurrence in perturbation trials just before target onset is an anomaly due to the inability to detect microsaccades during the blink owing to the blink-related eye movement. **b.** Rostral SC activity (same as middle column in Figure 24b) is shown in the bottom row for comparison. It is possible that the actual microsaccade rate increases during this period (suggested by the increase in rostral SC activity, particularly in monkey WM).

2.3.5 Alternate perturbations reveal mechanism mediating fixation-saccade network balance

We investigated whether the network-level effects observed above were specific to the perturbation method or extended to other ways of disrupting fixation. Specifically, we wanted to know if the perturbation acts by disrupting the flow of visual information and/or the motor act of fixating, both of which are affected when the animal blinks. It is also possible that the perturbation acts as a higher order cognitive cue informing the animal of impending target onset, thereby switching the network to a different state of preparation. In order to delineate these potential mechanisms, we also employed other forms of perturbations. First, we blanked out the central target during fixation at approximately the same times the blink would have occurred. This manipulation disrupts visual input to the fixation network while preserving fixation itself. Next, as a control perturbation to test the hypothesis that the cue may act in a top-down fashion, we presented a puff of air to the animal's ear during fixation. Figures 26a and 26b show the histogram of saccade latencies across trials and proportion of premature saccades across sessions for the target blank and ear puff perturbations, respectively. The target blank was similar to the blink perturbation in terms of the distribution of premature saccades – there were significantly more delay period errors compared to the control condition (Wilcoxon signed-rank test, $n=13$, $p<0.0001$). In contrast, however, the animals were not more likely to make errors following the ear puff compared to the control condition (Wilcoxon signed-rank test, $n=11$, $p>0.5$). Thus, it is unlikely the blink and target blank perturbations acted exclusively as a cognitive cue to the animal.

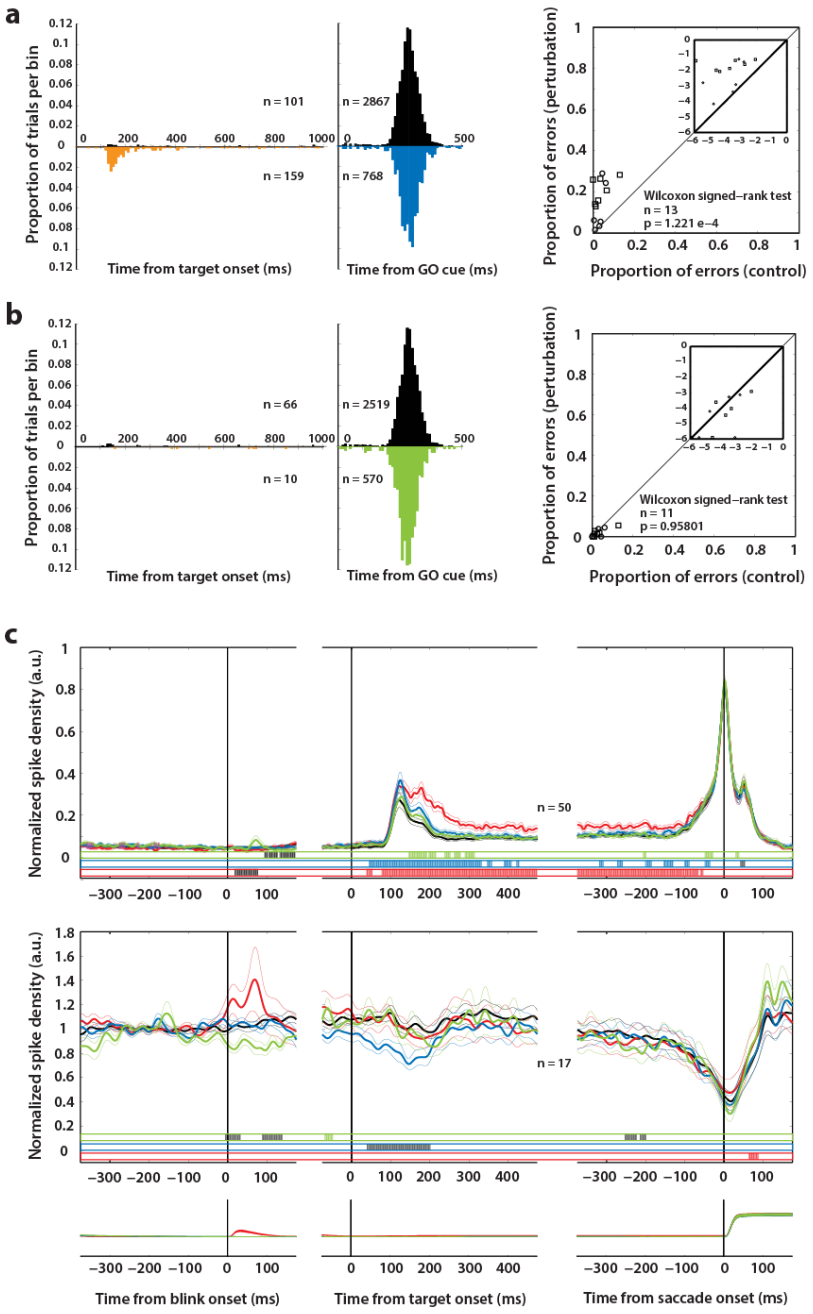

Figure 26. Behavior and population activity for alternative perturbations

a. Left. Histogram of saccade reaction times as in Figure 21b for control (top) and the target blank perturbation (bottom) trials. Right. Proportion of errors in the target blank perturbation condition plotted against the control condition. Each point represents a session; symbols correspond to different monkeys. The unity line is on the diagonal. **b.** Same as **a**, but for the ear puff perturbation. **c.** Population activity (mean +/- s.e.m.) of caudal SC (top) and rostral SC (middle) neurons for each of the four conditions (control – black, blink – red, target blank – blue, ear puff – green). The activity is plotted only for the matched subset of neurons for which we had data for all four conditions. The tick marks of a particular color indicate time points at which the corresponding condition had a higher activity relative to the control condition (Wilcoxon rank-sum test, $p<0.01$). Specific comparisons are enclosed by thin lines of the relevant color. The bottom row shows mean eye position traces.

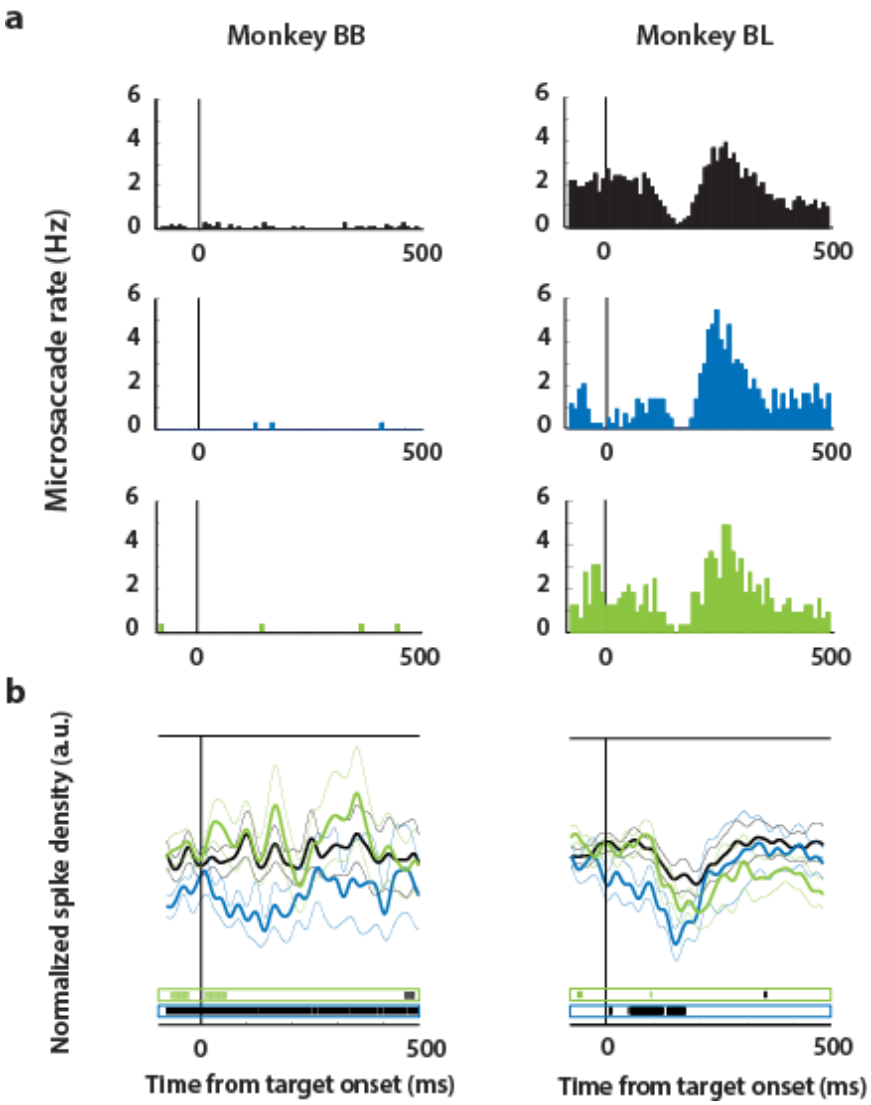

Figure 27. Microsaccade behaviour during alternative perturbations

a. Microsaccade rate as a function of time relative to target onset for two monkeys in control (top row; data in black), target blank (2nd row; data in blue) and ear puff (3rd row; data in green) conditions. Just like in the blink perturbation trials (Figure 25), monkey BB rarely made microsaccades at any time after acquiring fixation of the central fixation target. Monkey BL showed a characteristic microsaccade rate profile with transient inhibition of microsaccades following target onset. Since there is no blink to occlude microsaccades in the target blank trials (blue histogram), the observed reduction in microsaccade occurrence just before target onset could possibly be attributed to the reduction in rostral SC activity (bottom row). **b.** The bottom row plots rostral SC activity during the three conditions. The color configuration follows from above. The tick marks of a particular color indicate time points at which the corresponding condition had a higher activity relative to the control condition (Wilcoxon rank-sum test, $p<0.01$). Specific comparisons are enclosed by thin lines of the relevant color.

Figure 26c shows the population activity of caudal (top) and rostral SC (bottom) neurons from two animals, for these conditions, compared to the control and blink conditions for the same subset of neurons. Activity in caudal SC was enhanced and activity in rostral SC suppressed following the target blank perturbation (compare blue and black traces; Wilcoxon signed-rank test, $p<0.01$). The suppression in rostral SC seemed to be linked to a reduction in microsaccade rate in the one animal that made microsaccades in this condition (monkey BL; Figure 27a, middle row). There was minimal change to the network following the ear puff perturbation (green traces), although there was a slight increase in caudal SC activity during the late component of the visual response (Figure 26c, significance bars on bottom of each panel). Note that the red traces in Figure 26c are from blink perturbation trials recorded for the same subset of neurons. The effects were similar to those described above for the larger population of neurons. Together, these results suggest that the perturbation must necessarily disrupt the visual and/or motor aspects of fixation in order to produce the network-level effects observed in this study.

2.3.6 Activity during premature saccades is reflective of a caudal SC-driven shift in balance

Thus far, we have focused on the subset of perturbation trials that produced regular latency saccades in the delayed saccade task. Finally, we turn to the small proportion of blink perturbation trials that resulted in saccades erroneously triggered before the go cue ("gold" trials in Figure 21). Figure 28 shows the population activity of caudal SC (top row) neurons for blink perturbation trials in which a premature saccade was produced (gold traces), compared to trials with regular latency saccades in the perturbed (red) and control (black) conditions for the same subset of neurons. Note that the difference between regular latency saccades in control and blink-perturbed trials for this subset of caudal SC neurons is similar to the effect described earlier in Figure 22 – an enhancement of visuomovement activity following target onset in the perturbed condition, starting 94 ms after target onset (Figure 28, top row, middle column, red vs black traces; also see red significance bar below activity traces). Caudal SC activity was even higher on premature saccade trials, and separated earlier (54 ms after target onset; see gold significance bar below activity traces) from activity in the control condition (gold vs black traces). Importantly, perturbation trials with premature saccades differentiated from trials with regular latency saccades

(gold vs red traces) during the early part of visual epoch (shaded window), with the respective activity profiles separating 92 ms after target onset (see arrow in significance band above activity traces).

Rostral SC activity aligned on target onset was expectedly lower during premature saccades compared to when there was no early saccade following target onset in the control or blink-perturbed conditions (Figure 28, middle row, middle column, gold vs black and red traces). Intriguingly, compared to caudal SC, rostral SC activity for perturbation trials with premature saccades separated from trials with regular latency saccades much later, around 200 ms after target onset (see arrow in middle row). This observation is consistent with the interpretation that premature movements are primarily driven by an increase in visuomovement activity in caudal SC. It is important to highlight that visuomovement activity profiles during the visual epoch in caudal SC (Figure 28, top row, middle column, black, red, and gold trials) seem to fall on a continuum, suggesting that the stimulus-evoked response following the perturbation is graded and may sometimes be strong enough to trigger a movement. Drawing from this observation, we wish to re-emphasize the idea that sub-threshold (below the phenomenological threshold to evoke movements) neural changes can reveal the evolution of mechanisms leading to premature movements.

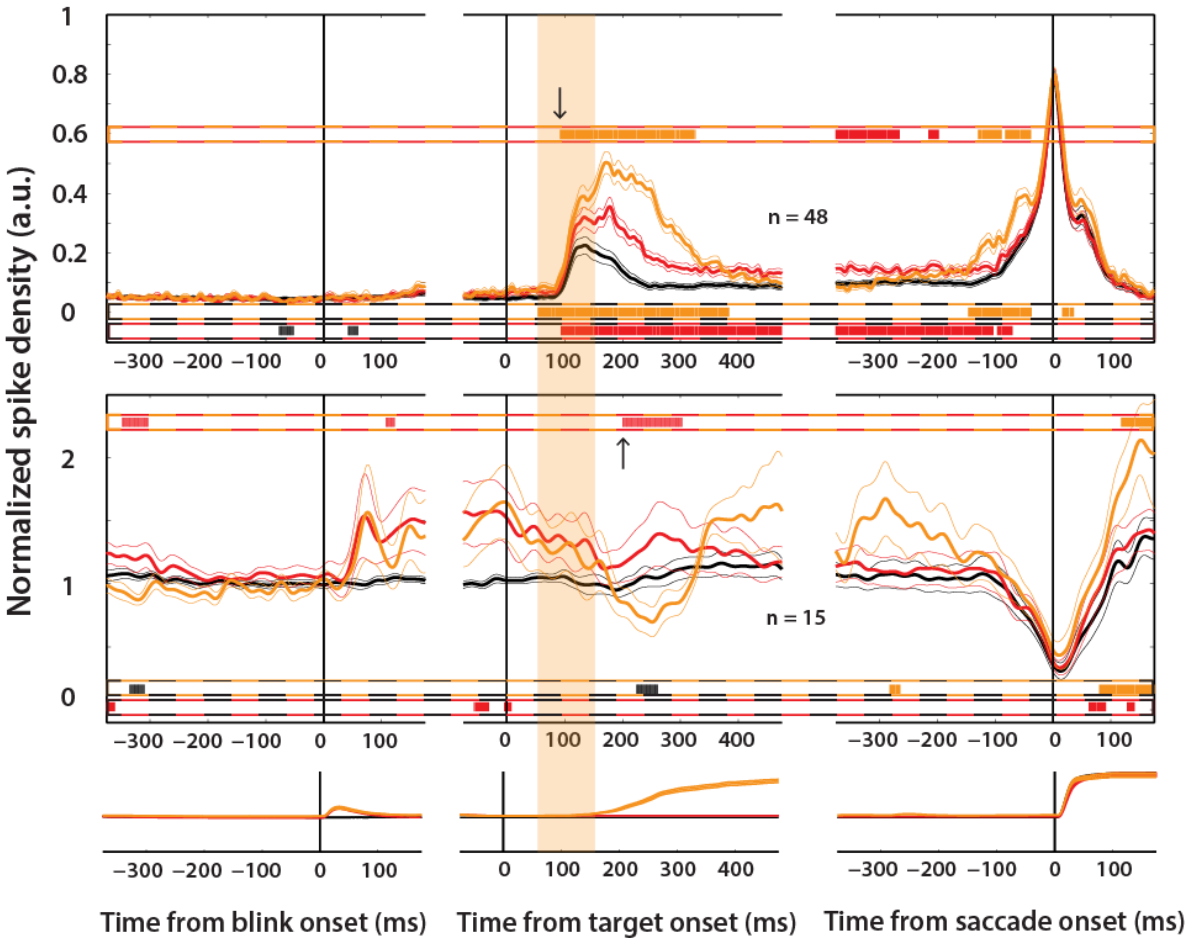

Figure 28. Population activity during premature saccades

Normalized population average activity (mean +/- s.e.m.) of caudal (top row) and rostral (middle row) SC neurons for control trials (black), blink perturbation trials with regular latency saccades (red), and blink perturbation trials with early saccades (gold). Note that for the gold trials, the activity in both middle and right panels reflect saccade occurrences, but are aligned with respect to different events (target onset and saccade onset, respectively). The two sets of significance bars below the spike density profiles, enclosed by dashed black and colored lines, indicate time points at which the condition corresponding to the respective color had higher (colored bars) or lower (black bars) activity relative to the control condition. The significance bars above the spike densities, enclosed by dashed red and gold lines, indicate time points at which activity in the gold trials was higher (gold bars) or lower (red bars) compared to the red trials; the arrows indicate first times at which they significantly separated from each other (Wilcoxon rank-sum test, $p<0.01$). The shaded vertical rectangle indicates the early visual epoch, 50-150 ms after target onset. The bottom row shows the mean +/- s.e.m. of the vectorial eye position in the three conditions.

2.4 DISCUSSION

We probed the functional interactions between movement-generating neurons and movement-suppressing neurons involved in the control of gaze by perturbing fixation behaviorally, with reflex blinks, while recording neural activity in the caudal and rostral portions of SC. We found that the perturbation produced profound changes in the neural activity even when the animal's reaction time was not compromised. It increased the intensity of the visual burst in visuomovement neurons in caudal SC and revealed novel sensorimotor properties of these saccade-generating neurons. Neurons which were more movement-like along the visuomovement continuum were more likely to be affected by the perturbation and tended to exhibit more visual-like properties. We also found that disruption of fixation had a complex effect on rostral SC. The tonic activity of rostral SC neurons exhibited minimal changes at the population level across subjects but showed dramatic differences between them. We attributed some of these differences to the occurrence of microsaccades. We considered alternative explanations for these observations and showed that other methods of perturbing fixation also had similar effects on the rostral-caudal SC network; however, it was important that the perturbation disrupted fixation either visually or by inducing an eye movement. A control perturbation with the potential to serve as a cognitive cue in advance of target onset, but that did not disrupt fixation, produced minimal effect on the network. Finally, we showed that perturbation-induced shifts are primarily driven by an increase in visuomovement activity in caudal SC which may sometimes trigger erroneous or premature saccades.

What do these results reveal about sensorimotor processing in the gaze control network? In the delayed saccade task, the monkeys were trained to suppress a saccade immediately following the onset of a target and hold their fixation until the GO cue. Under these conditions, visuomovement neurons in SC produce an initial burst of activity locked to target onset but is typically insufficient to trigger a saccade. Visuomovement neurons in SC project directly to the saccade burst generators in the brainstem (Rodgers et al., 2006), and an increase in their activity increases the drive to generate a saccade. Thus, the "visual response" is thought to be an embodied signature of the sensorimotor system registering the presence of the saccade target. In other words, it represents the latent intention to generate a saccade to the corresponding target. When fixation is disrupted, the target's representation is enhanced, indicating that the drive to look at the target increases. This interpretation is supported by the observation that on occasion, the enhancement is

sufficiently large to trigger a premature saccade (Figure 28). It is important to note that the perturbation occurred tens to hundreds of milliseconds before target onset, and hence, these effects reflect persistent, long-lasting changes to the balance between the fixation and saccade-generating networks – changes that do not necessarily have a behavioral signature. However, the enhancement of the visual response may reflect internal changes, e.g., greater visual attention and/or increased target salience. Although the interpretation of a flexible visual-movement continuum is based on a wealth of studies that have found little differences between visual and saccade-related responses of visuomovement neurons, it is important to mention recent studies that have shown that a systematic transformation occurs between the visual and motor responses, from representing target location in space to representing actual movement metrics (Sadeh et al., 2015; Sajad et al., 2015). Our results suggest a lack of strong functional distinction between the visual and movement bursts, but specifically with respect to initiating the movement. It is still possible that variation in the amplitude of the motor burst is correlated with variation in the executed saccade across trials. If so, the observed enhancement of the visual burst is consistent with the notion that the sensory response is not correlated with the actual movement.

Enhanced excitability of the saccade-generating network seems to occur through facilitation of the columnar dorso-ventral pathways within the SC or through its extra-collicular inputs. It does not seem to be directly related to the inability to maintain fixation as judged from the heterogeneous response of rostral SC fixation neurons. In one subject, the increased visuomovement activation was accompanied by a concomitant reduction in fixation-related activity, whereas in the others, it was associated with an increase in the firing of rostral SC neurons. Previous studies have linked such inter-subject differences in the state of SC activation to idiosyncracies in microsaccade behavior (Hafed et al., 2013). It is possible that subjects that for some reason generate few or no microsaccades have a strong functional link (or one that is apparent) between the fixation and saccade-generating networks. Other subjects that may have developed a microsaccade-based strategy for maintaining fixation may have a weaker reciprocal relationship between the two systems. In these subjects, behavioral disruption of fixation may increase activity across the SC map, and the resulting impulsive behavior may manifest as microsaccades or saccades alike. This latter interpretation is consistent with the hypothesis that microsaccades and saccades (and their associated neural substrates) form a functional continuum (Zuber and Stark, 1965).

In addition to its role in enabling fixation and fixational eye movements, it has been proposed that the rostral SC is involved in maintaining gaze-related goals of the animal (Krauzlis et al., 2004). In a multi-step gaze shift task, fixation neurons in rostral SC were shown to suppress their activity for the first step, and remained so until the end of the trial, turning on only when the complete sequence of movements was completed (Bergeron and Guitton, 2000). In our study, at least in one subject, it is possible that the disruption of fixation marked the initial step, causing a sustained suppression of rostral SC activity.

We considered several alternative mechanisms leading to the results in this study. The enhancement of the visual response is similar to that observed during PFC cooling (Koval et al., 2011). The PFC has been implicated extensively in the executive control of movements (Funahashi, 2001), including saccades, and PFC deactivation or damage results in an inability to inhibit premature responses driven by the stimulus (Pierrot-Deseilligny et al., 1991). These observations together suggest that the effect reported in this study may be due to an altered balance in the PFC-SC network. Could the perturbation have served as a cognitive cue to the animal about impending target onset, allowing top-down biasing of the SC network balance? We have seen that a couple of different perturbations produced the observed effects in this study. However, the control perturbation of a puff of air to the ear had a weak, if discernible, effect. We suggest that the best possible interpretation given these observations is that even if the perturbation is a top-down cue, it must act through the visual-oculomotor pathway in order to effectively disrupt fixation and create a deficit that resembles executive control disorders. Another possibility is that the perturbation acted as a 'startle cue', and produced a startle response. Unexpected startle cues have been shown to excite the sensorimotor pathway (Kumru and Valls-Sole, 2006) and accelerate the execution of planned movements (Castellote et al., 2007). Since the perturbation was only applied on a small percentage of trials, it may have produced a startle response, resulting in the disinhibition of visuomovement neurons. However, the fact that the puff of air to the ear only had a weak effect suggests that any startle-related changes must act through the visuomotor pathway. Regardless of which of the above mechanisms was at play, the effective result of the perturbations was an altered balance between movement generation and inhibition in the gaze control network.

Note that disruption of fixation occasionally resulted in premature eye movements (gold data in Figures 21 and 28). We propose that this mimics impulse control disorders in which the

ability to appropriately transition from fixation to saccade generation is compromised. Such deficits have been observed in neuropsychiatric disorders such as ADHD (Munoz et al., 2003) and schizophrenia (Braff and Geyer, 1990; Crawford et al., 2002), and movement control disorders like Parkinson's disease (Terao et al., 2011). Indeed, impulsivity and pathological gaze control is used as a common behavioral assay for executive dysfunction (Reuter and Kathmann, 2004). We propose that in order to gain a deeper understanding into the neural mechanisms causing impulsive movements, it is important to have access to internal changes that do not always manifest in overt behavior. Previous studies have attempted to study functional interactions in the gaze control network by using causal manipulations such as microstimulation or pharmacological inactivation. However, the understanding that can be gained with those manipulations is limited considering the fact that it is difficult to record activity in the affected neuronal populations. Moreover, the invasive nature of those manipulations makes it difficult to apply them to human experiments. The simple, non-invasive, behavioral perturbations we used here are readily accessible for both animal and human studies. Importantly, they allow us to examine latent network processes. Such perturbation approaches may prove to be an essential supplement to standard neurophysiological experiments enabling a complete understanding of neural mechanisms underlying various behaviors.

3.0 EXTRACTING SACCADES FROM EYE MOVEMENTS TRIGGERED BY REFLEX BLINKS

3.1 INTRODUCTION

The saccadic system has proven to be a powerful model to study the neural mechanisms underlying movement preparation. Numerous studies have probed nodes in the gaze control network involved in saccade programming in an attempt to delineate the time course of the preparatory process (Hanes and Schall, 1996; Dorris and Munoz, 1998), specifically, to understand when sensory and cognitive processes end and premotor processing begins (Kustov and Robinson, 1996; Thompson et al., 1996; Horwitz and Newsome, 1999; Gold and Shadlen, 2000; Dorris et al., 2007). This approach has generated a wealth of data about the neural correlates of saccade preparation during normal behaviour. However, in these studies, the relationship between preparatory activity and behaviour is inferred by correlating two processes (neural activity and movement generation) that are separated by hundreds of milliseconds. Thus, it is unclear whether evolving sub-threshold activity in the gaze control network also has an underlying behavioural correlate. Other studies have used causal approaches like microstimulation to gauge the presence and dynamics of motor signals in ongoing activity (Juan et al., 2004), but the interpretation of these results is confounded by the introduction of extraneous activity.

Recent work in our lab has used reflex blinks as a tool to obtain a behavioral readout of the (normally hidden) preparatory process. Appropriately timed reflex blinks reduce saccade reaction times significantly (Gandhi and Bonadonna, 2005), an observation that has been used to study the link between target selection and movement preparation in a visual search task (Katnani and Gandhi, 2013). Blinks remove inhibition on the saccadic system by shutting down the omnipause neurons (OPNs) in the brainstem – these neurons are tonically active when the eyes are stable and turn off during eye movements, thus gating them (Cohen and Henn, 1972). The OPNs turn off

during blinks due to the fact that the blink induces a slow, loopy, blink-related eye movement (BREM) (Rottach et al., 1998; Schultz et al., 2010)., and we believe the OPN inhibition allows a saccade to follow and overlap with the BREM (Gandhi and Bonadonna, 2005). We refer to such movements as blink-triggered saccades.

Figure 31a shows examples of BREMs evoked during fixation of a central target. Note that the excursion of the eye is significantly large (~5-10 degrees), although this can vary depending on the subject or initial eye position (Rottach et al., 1998). As a consequence, reduced latency saccades triggered by reflex blinks are contaminated by the presence of the BREM. Figure 31b shows examples of normal saccades and blink-triggered gaze shifts – note that the spatial and temporal profiles of the two sets of movements are different. The initial phase of blink-triggered movements typically overlaps with the initial phase of BREMs (see overlaid purple traces in Figure 31b). In order to use blinks to study temporal aspects of the preparatory process, it is important to know when exactly during the course of the blink-triggered movement the actual saccade is initiated.

Previous studies have shown that when a reflex blink is evoked just around saccade onset, the resultant eye movement is slower and does not follow typical saccade kinematics. Such "blink-perturbed" saccades cannot be modeled as a linear combination of the BREM and a control saccade directed to the same stimulus (Goossens and Van Opstal, 2000a). Moreover, activity in a subset of saccade-generating neurons in the superior colliculus (SC) is suppressed during blink-perturbed movements (Goossens and Van Opstal, 2000b). These studies used a velocity threshold-based criterion to determine the saccadic component of the complex blink-perturbed movement. Here, we show that blink-triggered movements can in fact be modeled as a linear combination of a typical BREM and a typical saccade, crucially, with an imposed delay between the two components. Saccades reconstructed from the blink-triggered movement using this approach were largely similar to normal saccades in their spatial and temporal dynamics. We further show that saccade-related bursts in SC for the recovered saccade closely resemble those for normal saccades. We compare our results against the previously used velocity threshold-based criterion. Our results therefore support the notion that blink perturbations, particularly those that occur well before saccade onset, if properly accounted for, offer a non-invasive tool to probe the behavioral and neural signatures of sensory-to-motor transformations.

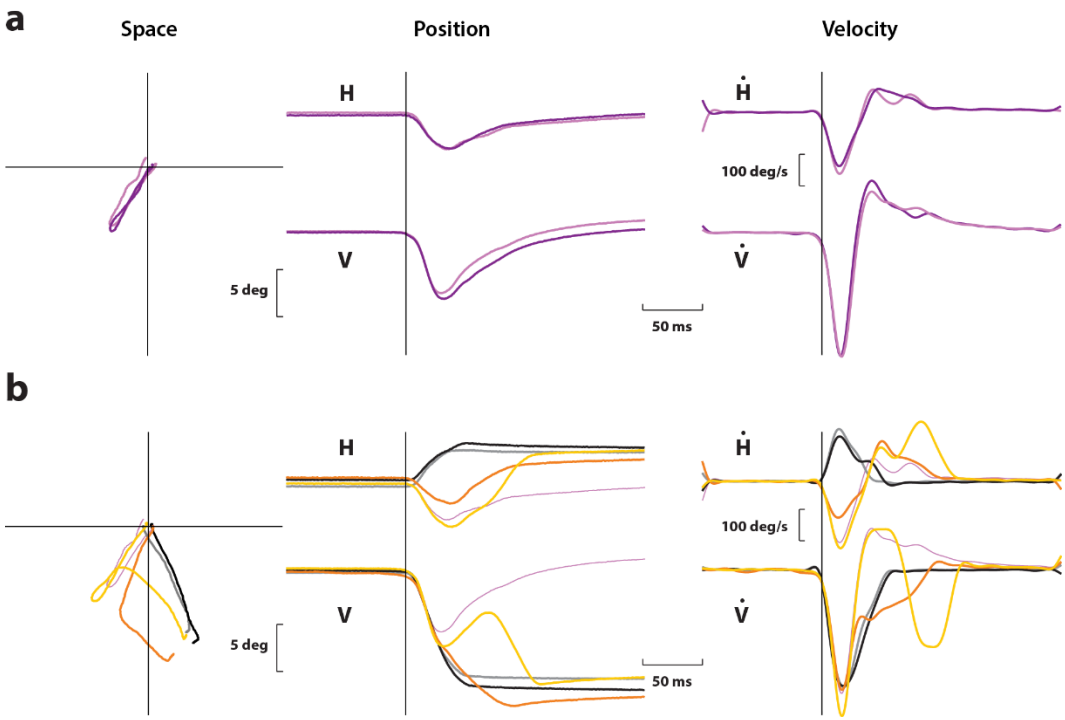

Figure 31. Blink-related eye movements and blink-triggered gaze shifts

a. Example spatial (left), position (middle), and velocity (right) profiles of two BREMs evoked during fixation. The two rows in the middle and right columns show horizontal and vertical signals, respectively. **b**. Examples of two normal saccades (gray to black) and two blink-triggered gaze shifts (orange to yellow) made to the same target. The layout is the same as in **a**. Note the heterogeneity in the profiles of blink-triggered movements. One of the BREMs (thin purple traces) from **a** is overlaid for comparison. Note the similarity between the initial phases of blink-triggered movements and the BREM.

3.2 METHODS

3.2.1 General and surgical procedures

All experimental and surgical procedures were approved by the Institutional Animal Care and Use Committee at the University of Pittsburgh and were in compliance with the US Public Health Service policy on the humane care and use of laboratory animals. We used two adult rhesus monkeys (*Macaca mulatta*, 1 male and 1 female, ages 8 and 10, respectively) for our experiments. Under isoflurane anesthesia, a craniotomy that allowed access to the SC was performed and a recording chamber was secured to the skull over the craniotomy. In addition, posts for head restraint and scleral search coils to track gaze shifts were implanted. Post-recovery, the animal was trained to perform standard eye movement tasks for a liquid reward.

3.2.2 Visual stimuli and behavior

Visual stimuli were displayed by back-projection onto a hemispherical dome. Stimuli were white squares on a dark grey background, 4x4 pixels in size and subtended approximately 0.5° of visual angle. Eye position was recorded using the scleral search coil technique, sampled at 1 kHz. Stimulus presentation and the animal's behavior were under real-time control with a LabVIEW-based controller interface (Bryant and Gandhi, 2005). After initial training and acclimatization, the monkeys were trained to perform a delayed saccade task. The subject was required to initiate the trial by acquiring fixation on a central fixation target. Next, a target appeared in the periphery but the fixation point remained illuminated for a variable 500-1200 ms, and the animal was required to delay saccade onset until the fixation point was extinguished (GO cue). Trials in which fixation was broken before peripheral target onset were removed from further analyses. The animals performed the task correctly on >95% of the trials.

3.2.3 Induction of reflex blinks

On a small percentage of trials (~15-20%), we delivered an air puff to the animal's eye to invoke the trigeminal blink reflex. Compressed air was fed through a pressure valve and air flow was monitored with a flow meter. To record blinks, we taped a small Teflon-coated stainless steel coil (similar to the ones used for eye tracking, but smaller in coil diameter) to the top of the eyelid. The air pressure was titrated during each session to evoke a single blink. Trials in which the animal blinked excessively or did not blink were aborted and/or excluded from further analyses. To obtain blink-triggered movements to the peripheral target, we sought to evoke blinks 100-250 ms after the GO cue, during the early phase of the typical saccade reaction time. In our experimental setup, blink onset occurs approximately 150 ms after the air puff reaches the eye. Thus, air puffs were administered 50 ms before to 100 ms after the GO cue. Trials in which a saccade did not accompany such blinks (i.e., where the gaze did not end up at the target) were removed from further analysis. To obtain BREMs without an accompanying saccade, air puff delivery was timed to evoke blinks during fixation of the central target, 400-100 ms before the onset of the peripheral target. The window constraints for gaze were relaxed for a period of 200-500 ms following delivery of the air puff to ensure that the excursion of the BREM did not lead to an aborted trial.

3.2.4 Electrophysiology

During each recording session, a tungsten microelectrode was lowered into the SC chamber using a hydraulic microdrive. Neural activity was amplified and band-pass filtered between 200 Hz and 5 kHz and fed to a digital oscilloscope for visualization and spike discrimination. A window discriminator was used to threshold and trigger spikes online, and the corresponding spike times were recorded. The location of the electrode in the SC was confirmed by the presence of visual and movement-related activity as well as the ability to evoke fixed vector saccadic eye movements at low stimulation currents (20-40 μA, 400 Hz, 100 ms). Before beginning data collection for a given neuron, its response field was roughly estimated. During data collection, the saccade target was placed either in the neuron's response field or at the diametrically opposite location (reflected across both axes) in a randomly interleaved manner.

3.3 DATA ANALYSIS

3.3.1 Data analysis and pre-processing

Data were analyzed using a combination of in-house software and Matlab. Eye position signals were smoothed with a phase-neutral filter and differentiated to obtain velocity traces. Normal saccades, BREMs, and blink-triggered eye movements were detected using standard onset and offset velocity criteria (50 deg/s and 30 deg/s, respectively). Onsets and offsets were detected separately for horizontal and vertical components of the movements and the minimum (maximum) of the two values was taken to be the actual onset (offset).

Raw spike density waveforms were computed for each neuron and each trial by convolving the spike trains with a Gaussian kernel (width = 4 ms). For a given neuron and target location, spike densities were averaged across trials after aligning to saccade onset. We also normalized the trial-averaged spike density of each neuron to enable meaningful averaging across the population. The activity of each neuron was normalized by its peak firing rate during normal saccades.

3.3.2 Modeling of blink-triggered movements

Our objective in this study was to test whether a linear model that took into account saccade and BREM dynamics was sufficient to account for the heterogeneity of blink-triggered movements. If this was indeed true, the saccade could be extracted from linear decomposition of the blink-triggered movement. Previous studies had found that blink-perturbed movements could not be explained as a simple linear combination of saccades and BREMs (Goossens and Van Opstal, 2000a). Hence, we created a modified linear model – linear combination with delay. We chose velocity space to model the movements because of the sharper array of features (e.g., number of peaks) in the velocity signal compared to the position signal that could potentially be exploited to fit the movements. The model is described by the following equation:

$$\dot{\boldsymbol{M}}(t) = \dot{\boldsymbol{b}}_j(t) + \dot{\boldsymbol{s}}_k(t - \Delta)$$

where $\boldsymbol{M}(t)$, $\boldsymbol{b}(t)$, and $\boldsymbol{s}(t)$ represent the position signals of the simulated blink-triggered movement, BREM, and normal saccade, respectively, and the overdot denotes the derivative of

those signals, representing velocity. Note that each of the terms is a two-element vector containing values from horizontal and vertical channels. For each of the three movements, we included additional 20 ms snippets on either side of the onset and offset times determined by velocity criteria.

Next, we minimized Euclidean distance in velocity space to fit the simulated movements to recorded blink-triggered movements. The free parameters of the model are j, k, and Δ, which are the indices of individual BREMs, saccades, and a time delay. For each blink-triggered movement $\boldsymbol{m}_i(t)$, we determined the optimal parameters as

$$(j_{opt}, k_{opt}, \Delta_{\mathrm{opt}}) = argmin \left|\left|\dot{\boldsymbol{M}}(t) - \dot{\boldsymbol{m}}_i(t)\right|\right|^2$$

The parameter space was chosen to be computationally efficient and spanning reasonable values. Thus, we randomly chose 100 trials to provide the saccade profiles, 50 trials to provide BREM profiles, and scanned integer time delays in the interval $[-100, 100]$. Negative values for Δ_{opt} indicate that saccade onset precedes BREM onset in the blink-triggered movement, while positive values indicate that saccade onset follows BREM onset. This approach is illustrated in Figure 32a.

3.3.3 Detection of back-thresholded saccades

For comparison with the linear combination with delay model, we also detected saccade onset in blink-triggered movements using a previously used backwards velocity threshold method (Gandhi and Bonadonna, 2005). In order to do this, we detected the peak in the velocity profile of the blink-triggered movement and marched backwards in time until the standard onset velocity criterion (50 deg/s) was crossed for at least 5 consecutive time points. The idea behind this approach is that the peak velocity of the movement should occur during the saccade component, and going backwards in time until the threshold is crossed should isolate the saccade alone. In several sessions, peak velocity was attained during the initial phase of the blink-triggered movement, which was clearly contributed by the BREM. Hence, we applied the backwards threshold method starting from the second significant peak for these movements. This approach is depicted in Figure 32b.

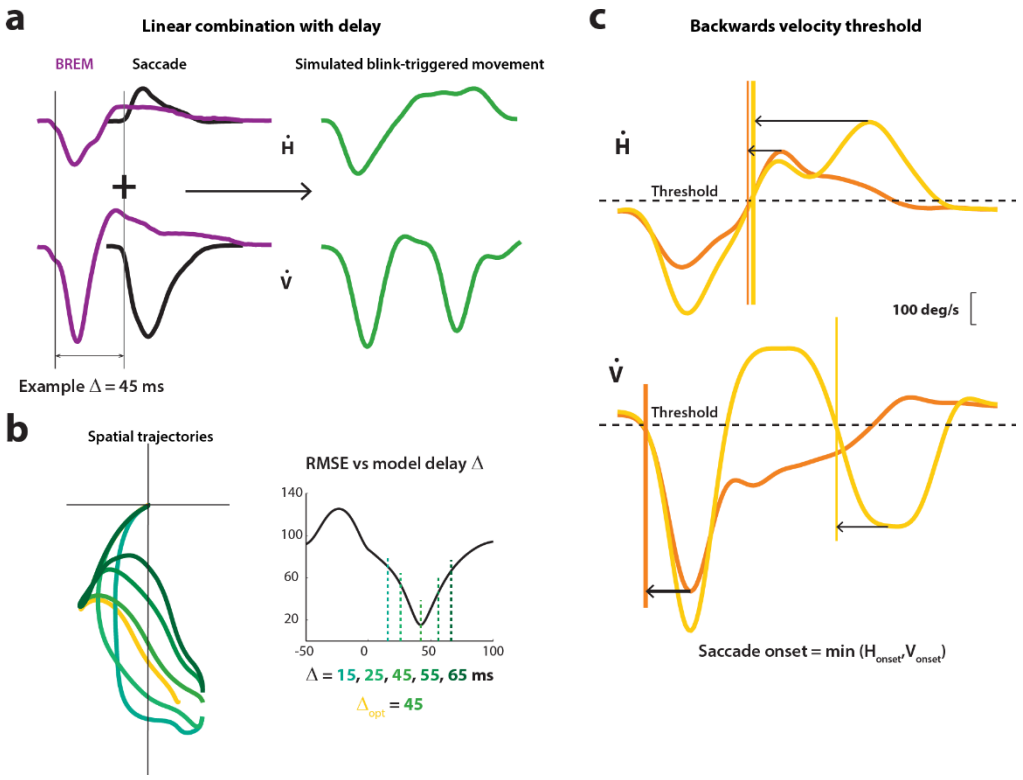

Figure 32. Saccade extraction algorithms

a. Schematic of the linear combination approach. The velocity profiles of BREMs (purple traces) and normal saccades (black traces) are shifted in time (e.g, by 45 ms) and added to give a simulated blink-triggered movement (green traces). **b.** Spatial trajectories of an example blink-triggered movement (yellow trace) and five simulated movements (shades of green) are shown on the left. On the right is a plot of the root mean-squared error between various simulated movement profiles and the actual profile as a function of simulated time shift Δ. The value of the time shift that provides the best fit to the actual blink-triggered movement (in this example, the best fit occurs at delay = 45 ms) is considered to be the starting point of the saccade in the combined movement. **c**. Schematic of the backwards threshold approach. Two example blink-triggered movement traces (yellow and orange) from Figure 31 are shown. The time at which velocity crosses a pre-determined threshold (horizontal dotted line), computed backwards from the last peak, is considered as the onset time for that component. The minimum of horizontal and vertical onset times is taken to be the time of saccade onset.

3.3.4 Other analyses and statistical tests

For the analyses that involve studying the effect on model performance of an external parameter (e.g., direction or optimal delay), we binned trials from all sessions according to that parameter, and computed the average mean-squared error (MSE) across trials for each bin. Significant trends were identified by comparing these to the null distribution (uniform distribution/no trend) generated by a bootstrap approach with appropriate confidence intervals that took into account the

number of trials available in a given bin. Where applicable, we used the Wilcoxon-rank-sum test for a two-way comparison of distribution medians.

In addition to MSE, we also used Spearman's rank correlation to compare the efficiency of different approaches in generating velocity profiles that closely match the control profiles. A point-by-point correlation of two profiles ignores magnitude differences (that is represented in the MSE) and tests whether their shapes are similar. Scaled profiles would result in a correlation of 1.

We also computed a suppression index for individual neurons as the relative change in the activity for extracted saccade trials with respect to control trials. We computed the index using the average activity in a 40 ms window around saccade onset as, $SI = \frac{r_{-20:20}^{control} - r_{-20:20}^{extracted}}{r_{-20:20}^{control}}$, where $r_{window}^{condition}$ is the average activity in $window$ around saccade onset for $condition$ trials.

3.4 RESULTS

In order to extract saccades from blink-triggered movements, we needed three sets of data. In animals performing the delayed saccade task, during each session, we recorded 1) a set of normal saccades to two or more targets (typically more than 100 per session), 2) a set of blink-triggered gaze shifts to those targets by inducing reflex blinks after the go cue (around 100 trials per session), 3) blink-related eye movement (BREM) profiles by evoking blinks during initial fixation, before peripheral target onset (around 50 per session). We had 100 session-target pairs in all, and analyzed data for each session-target pair separately.

First, we simulated a range of movement profiles by linear summation of randomly chosen saccade and BREM velocity profiles at all possible time shifts relative to each other. Then, for each real blink-triggered movement, we determined the simulated profile that best fit it (see Methods), and the associated triad of saccade, BREM and time shift parametrized the decomposition of the blink-triggered movement. Note that the optimal time shift effectively determined the time of saccade onset within the movement. Figure 33a shows the average blink-triggered movement velocity profiles for one session-target pair (orange traces) and the corresponding best-fitting simulated movement profile average (green traces). Next, we extracted

the underlying saccade by subtracting the respective BREM profile from each blink-triggered movement at the appropriate time shift. The dynamics of the extracted saccades closely resembled normal saccades, as judged from the shape of their velocity profiles (Figure 33b, compare blue vs black traces). For comparison, we also computed an alternative saccade onset time using a backwards velocity threshold criterion (see Methods) that has been used previously (Gandhi and Bonadonna, 2005). Note that the dynamics of the saccade estimated using this criterion (Figure 33b, red traces) deviated significantly from that of normal saccades or those extracted using the aforementioned method.

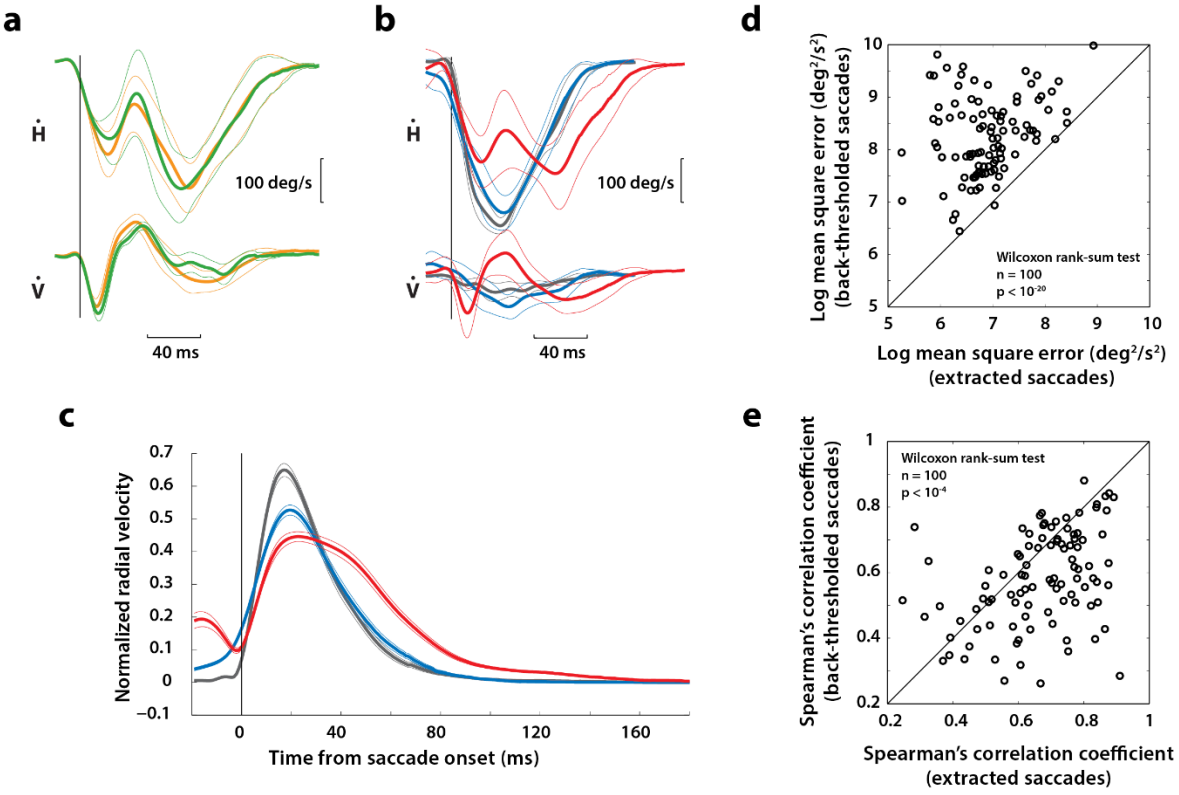

Figure 33. Comparison of model performance – behaviour

a. Average horizontal and vertical velocity profiles of actual (orange traces) and simulated best-fitting (green traces) blink-triggered movements for an example session-target pair. The thin lines are the standard error bounds across trials. **b**. Average horizontal and vertical velocity profiles of normal saccades (black), extracted saccades (blue), and back-thresholded saccades (red) for the same session-target pair as in **a**. **c**. Average normalized radial velocity profiles across all sessions for the three saccade types shown in **b**. **d**. Log-log plot of the mean-squared error between extracted and normal saccade velocity profiles (abscissa) plotted against back-thresholded and normal velocity profiles (red bars) across session-target pairs. Each point is from a session-target pair. **e**. Spearman's rank correlation between extracted and normal saccade velocity profiles plotted against the correlation between back-thresholded and normal velocity profiles.

Figure 33c shows the average velocity profiles across all session-target pairs for normal, extracted and back-thresholded saccades. Qualitatively, the extracted saccades overlap substantially with normal saccades. Some of the discrepancies could be because the dynamics of the motor command, and therefore the saccade, triggered in response to a blink may be slightly attenuated (for more, see below). Moreover, the velocity profiles of extracted saccades were much closer to normal saccades than those of back-thresholded saccades. To quantify this observation, we computed the mean square error (MSE, equivalently, Euclidean distance in velocity phase space) between normal and extracted saccades, as well as between normal and back-thresholded saccades, for each session-target pair. Figure 33d shows the MSEs for the two cases plotted against each other (blue – normal vs extracted, red – normal vs back-thresholded). The MSEs were significantly lower for extracted saccades compared to back-thresholded saccades (Wilcoxon rank-sum test, n=100, $p < 10^{-20}$). Note that the MSE reflects differences between two profiles regardless of their shape or where the deviations occur. To further evaluate whether the shape of the extracted saccade profiles resembled that of normal saccades, we computed the correlation between the velocity traces for each pair considered above. Two traces with similar shapes (e.g., scaled versions of each other) should have a high correlation coefficient compared to traces of dissimilar shapes. Figure 33e shows a scatter plot of the session-wise correlation values (Spearman's correlation coefficient) for the two pairs of traces (normal vs extracted, mean +/- s.d. = 0.71 +/- 0.14; normal vs back-thresholded, mean +/- s.d. = 0.6 +/- 0.13). The correlations were significantly higher for extracted saccades (Wilcoxon rank-sum test, n=100, $p < 10^{-4}$). Taken together, these results suggest that extraction of saccades from blink-triggered movements by linear decomposition is superior to the currently used method of determining the saccadic component and is likely more veridical in its estimates of saccade onset and dynamics.

Since the BREMs are stereotypical in a given animal and always excurse in roughly the same direction from the point of fixation, it is reasonable to expect an approach like ours to have differential efficacy depending on direction of the intended saccade. We therefore wanted to verify whether the performance of the algorithm was equally good for blink-triggered movements to all directions. We binned trials across all sessions into different directions based on the direction of the actual movement (i.e., final eye position relative to initial eye position) and computed the average MSE in each bin. This is shown in Figure 34a (thick blue trace; thin traces represent +/- s.e.m.). For comparison with the null (uniform) profile, we computed the bootstrapped version of

this trace by randomly assigning trials to direction bins (thick cyan trace; thin traces represent +/- 95% confidence intervals). The MSEs were largely uniform across directions, expect for a slight increase for movements to the bottom right quadrant, and a slight decrease for movements to the bottom left quadrant. Intriguingly, the BREMs were always directed to the bottom left quadrant for all our subjects. We also performed this analysis for the backwards threshold approach, as shown in Figure 34b (red curves; magenta curves show the bootstrapped distribution). The MSEs were more asymmetric and dependent on direction for back-thresholded saccades, suggesting that a simple threshold criterion may ignore the complex manner in which BREMs and saccades may interact dependent on direction.

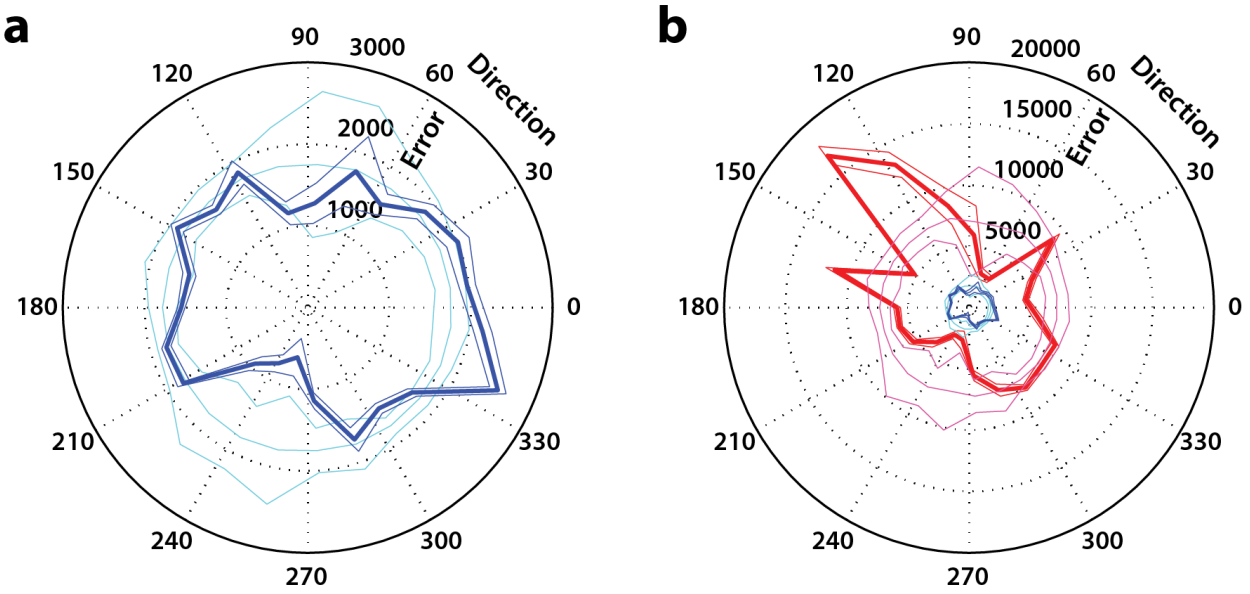

Figure 34. Robustness of model performance to saccade direction

a. Polar plot of the mean-squared error between individual extracted saccades and the average control saccade as a function of movement direction (blue traces). The thick blue trace is the average error across trials in a given direction bin, and the thin traces are +/- s.e.m. The cyan traces represent the average bootstrapped error and +/- 95% CI, and provide the null distribution in a given direction for comparison. **b.** Same as **a**, but for back-thresholded saccades (red traces), and the corresponding bootstrapped error distribution (magenta traces). The plot from **a** is also shown (blue traces near the centre) for comparison of magnitude.

It is pertinent to remark here that the method of linear combination is certain to provide an optimal decomposition of a blink-triggered movement, since any set of simulated movements is going to contain one that fits the actual movement best. Although the results presented above demonstrate the method's utility in extracting saccades with near-normal velocity profiles, its use in studies of the neural mechanisms of movement preparation requires evidence that saccade-

related neural activity is not unduly affected by the re-computed onset times. In order to determine whether this was true, we compared motor bursts of 51 neurons in the superior colliculus (SC) for normal saccades, and both extracted and back-thresholded saccades from blink-triggered movements. We only included session-target pairs in which the target was in the response field of the recorded SC neuron. Figure 35a shows the trial-averaged spike density of one neuron, aligned on respective saccade onsets, for each of the three cases. Note the remarkable similarity between the burst profiles for normal and extracted saccades (blue vs black traces). Figure 35b shows the average population activity for the three cases. To quantify the similarity of the bursts, we performed the correlation analysis presented earlier for the velocity profiles in Figure 33e. The distribution of correlation coefficients (Spearman's rank correlation) for the two pairs of comparisons plotted against each other is shown in Figure 35c. Correlations were significantly higher for the extracted saccade bursts (Wilcoxon rank-sum test, n=51, $p < 10^{-6}$), suggesting that linear decomposition with delay largely preserves the saccade-related burst compared to back-thresholding.

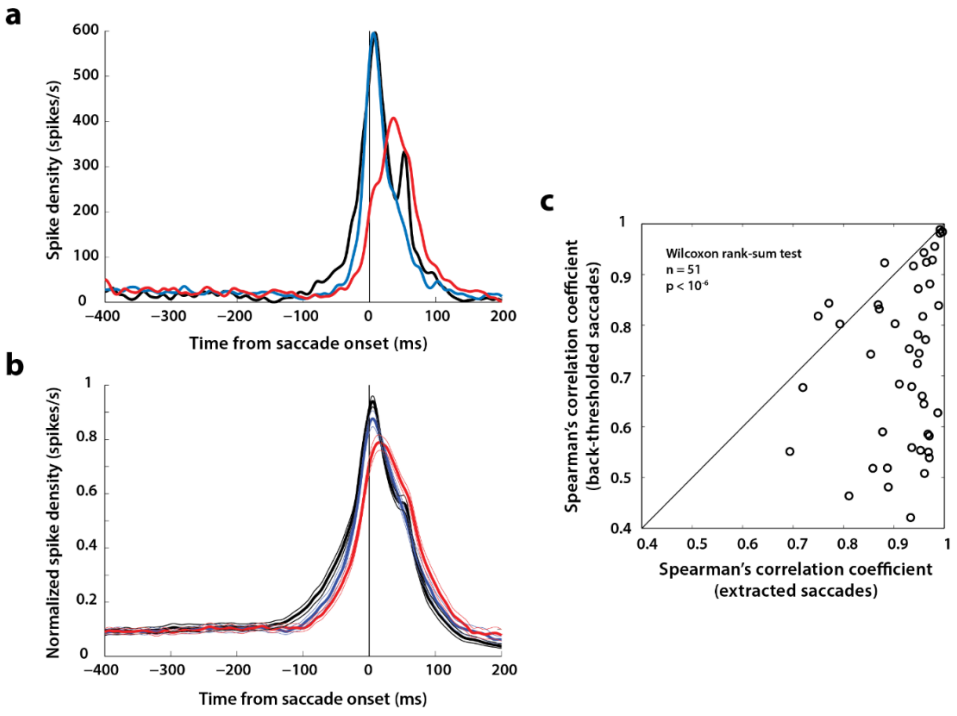

Figure 35. Comparison of model performance – neural activity

a. Saccade-related burst for one example SC neuron for normal (black), extracted (blue), and back-thresholded (red) saccades. **b**. Population-averaged motor burst in SC for the three types of saccades in **a**. **c**. Spearman's rank correlation between extracted and normal motor bursts plotted against the correlation between back-thresholded and normal motor bursts. Each point corresponds to one neuron.

Previous studies have observed strong attenuation in the saccadic burst in a few SC neurons (Goossens and Van Opstal, 2000b), but we did not. To resolve this discrepancy, we decided to take a closer look at the timing of the saccades extracted using our approach relative to the occurrence of the blink. Figure 36a shows the distribution of extracted saccade onset times relative to overall movement onset across all combined blink-saccade movements. In other words, this is the distribution of optimal time shifts at which linearly combining a BREM and saccade would produce each blink-triggered movement. Saccade onset followed BREM onset in the majority of trials. Observe that saccades can begin as late as 80 ms into the movement – an eon in the time scale of sensorimotor integration - highlighting the importance of precisely determining saccade onset time for studying movement preparation using the blink approach. Intriguingly, the distribution of optimal delays for saccade onset was bimodal, with an early peak around 10 ms and a later peak around 40 ms after movement onset. We thought this might provide a clue to the issue of suppression. Note that when saccade onset follows the blink by as early as 0-20 ms, the blink is occurring at a time when the command to generate a saccade has already been issued (accounting for efferent delays), and is likely to affect the burst when it is approaching peak value. If so, any suppression in these cases will alter the dynamics of the underlying programmed saccade and affect the ability of the algorithm to extract a control-like saccade.

We first tested whether the performance of the linear combination with delay approach was a function of the optimal delay identified by the algorithm (Figure 36b). The MSE was indeed higher for optimal delays between 0-20 ms (blue traces), indicated by the increase above the bootstrapped baseline (cyan traces). We then computed the suppression index for each neuron as a function of optimal delay bin. The average index across the population is shown in Figure 36c (thick trace; thin traces represent +/- s.e.m.). Note that the suppression is highest in the 0-20 ms window. Indeed, neuronal suppression was a reliable predictor of the MSE, as seen in the correlation of the two variables across bins (Figure 36d, $R^2 = 0.36$, $p < 0.01$). However, the magnitude of suppression was still lower (15% at its maximum in Figure 36c) than previously observed (Goossens and Van Opstal, 2000b). It is important to note that the suppression values so far were calculated based on aligning activity to extracted saccade onset, which seeks to maximize the similarity of the underlying saccade to a control movement (and may thus minimizes differences at the neural level as well), whereas, in previous studies, activity was aligned to saccade onset estimated using the backwards velocity threshold criterion. Suppression magnitudes were

indeed higher for back-thresholded saccades, peaking at 35% at the population level (Figure 36e). Finally, we wanted to verify whether individual neurons exhibited suppression that resembled the strong attenuation observed in the previous study. Figure 36f shows the average saccade-aligned burst profiles for five individual neurons for control (black traces) and back-thresholded saccades (red traces), in trials that fall within the 0-30 ms window of optimal delays estimated using the linear combination approach. In all five examples, a strong reduction/dip in the activity is evident during the peri-saccadic period. Thus, it is possible that the strong suppression observed in the previous study in a handful of neurons may be a result of a combination of factors, including timing of the saccade relative to the blink.

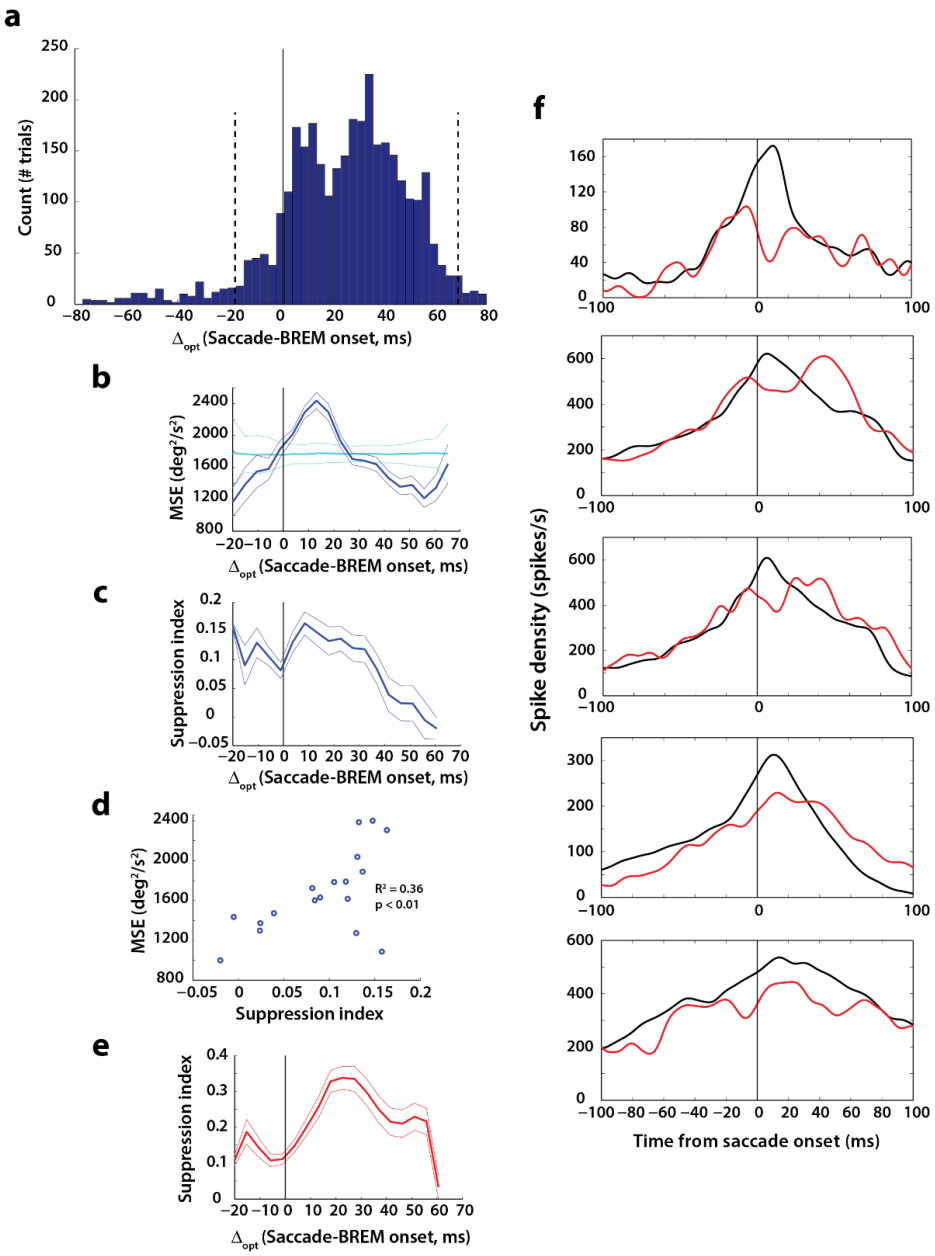

Figure 36. Time course of extracted saccades and activity suppression

a. Distribution of extracted saccade onset times relative to overall movement onset (delay, Δ_{opt}) for all trials. Based on the method, saccade onset occurred after movement onset (or, equivalently, BREM onset) in the majority of trials. The vertical dotted lines represent the 95th percentile of the distribution and the analyses were restricted to trials in this window of delays. **b**. Average mean-squared error between extracted and control saccade velocity profiles as a function of optimal delay (thick blue trace) and +/- s.e.m (thin blue traces). The cyan traces show the bootstrapped null distribution and +/- 95% CI for a test of uniformity. **c.** Neural suppression index (relative suppression for extracted saccades compared to control saccades) as a function of optimal delay. **d.** Correlation between mean-squared error and suppression index (points in **b** vs **c**). Each point is from a bin of optimal delays. **e.** Neural suppression index for back-thresholded saccades as a function of optimal delay, for better comparison with previous results. Note that relative suppression is higher when aligned with respect to the back-thresholded saccade, compared to that in **c**. **f.** Examples of suppression in five different neurons. Black traces show the trial-averaged activity aligned to saccade onset in control trials, and red traces depict the activity aligned to back-thresholded saccades, averaged across trials with optimal delays where suppression is highest (10-30 ms from **e**).

3.5 DISCUSSION

We have shown that blink-triggered movements can be characterized as a linear combination of BREMs and saccades with an arbitrary delay that can be estimated from the movement. Further, we saw that the dynamics of saccades extracted using this method as well as the associated motor bursts in SC are similar to those of normal saccades. A previous study of blink-perturbed saccades reported that they could not be explained by linear superposition of saccades and BREMs (Goossens and Van Opstal, 2000a). A follow-up study that analyzed SC activity during blink-perturbed saccades reported that the intensity of the saccade-related burst is attenuated for these movements while the relationship between spike count and saccade amplitude is preserved (Goossens and Van Opstal, 2000b). Thus, these studies concluded that blinks and saccades interact non-linearly at the level of the premotor circuitry in the saccadic system. The discrepancies between our results and the previous studies could be due to several factors. First, in the behavioural study, the possibility of a staggered, time-shifted linear superposition was not systematically explored (Goossens and Van Opstal, 2000a). As we show here, BREM onset and saccade onset can be offset by several tens of milliseconds, and taking this time shift into account can increase the explanatory power of the linear analysis significantly.

Second, we show that the saccade-related burst is largely preserved for the extracted saccade. While this is consistent with the previous observation that spike counts and saccade amplitude are correlated (if firing rates are preserved, so must be spike counts), it is inconsistent with the observed attenuation of the SC motor burst (Goossens and Van Opstal, 2000b). We believe this discrepancy could be due to the fact that our analyses included a large portion of blink-*triggered* saccades, i.e., movements in which blink onset preceded saccade onset by several tens of milliseconds, whereas previous studies focused largely on co-occurring blinks and saccades. Since our main focus was to use the reflex blink as a tool to probe underlying behavioral and neural processes, we induced blinks early, and thus obtained many trials where the blink likely triggered neural processes that resulted in a saccade much later. In the previous studies, blinks were timed to occur during typical saccade reaction times, likely resulting in instances where the saccadic burst had commenced and was perturbed by the blink. This notion is consistent with the observation that the suppression on blink-perturbed trials, especially for saccade onset times computed using previous approaches, is maximal around 20 ms after the onset of the BREM, which

is when the peak of the burst would have occurred if the final motor command was issued at the time of BREM (Figure 36e). In other words, we think the likelihood of observing suppression is higher when the blink interferes with the motor burst around the time of its peak. These claims are further bolstered by the observation that the algorithm fails to extract saccades resembling control saccades for this window of optimal delays (higher MSE, Figure 36b), possibly due to altered kinematics as a result of the attenuated burst peak, and the strong suppression of the peak in example individual neurons for trials that fall in this window of delays (Figure 36f). For true blink-triggered saccades, i.e., for longer optimal delays between blink and saccade onset, suppression of low-frequency preparatory activity, if any, may have enough time to recover to produce a full-fledged saccadic burst. The dynamics of saccades extracted from such blink-triggered movements do not seem to be very different from those of normal saccades; thus it is not surprising that SC activity for the two types of saccades are similar. It should be noted here that, even disregarding the specific focus on the time course of the saccade relative to the blink, the intensity of the motor burst was attenuated when we used the velocity-based threshold criterion (Figure 35b).

Another possible explanation for the discrepancy in the extent of attenuation is differences in the air puff stimuli used to induce reflex blinks. The observed short latency attenuation of saccade-related activity of a small subset of neurons in the previous study was locked to the time of air puff delivery, and as such, may have been linked to the strength of the air puff. It is entirely possible that the air puffs we used were weaker, and combined with early delivery times (when preparatory activity is lower), precluded the strong attenuation described previously. It should further be noted that the previous study observed strong suppression in only a small percentage of the overall population, and we observed a similar strong effect in several individual neurons (some of which are shown in Figure 36f). The primary purpose of this study was to find a way to precisely determine saccade onset from a blink-triggered movement, and we think a systematic analysis of these multiple factors and a thorough re-verification of previous results is beyond its scope.

The fact that BREMs and saccades can be linearly combined to produce the observed movement suggests that the underlying neural processes producing a BREM and saccade are independent. Although the neural correlates of a BREM are unclear, we have previously shown that the BREM itself does not affect activity in caudal SC (Jagadisan and Gandhi, 2016). However, the BREM has been shown to be associated with a pause in the activity of the omnipause neurons

(OPNs) in the brainstem (Schultz et al., 2010), which we think allows activity in premotor structures such as SC to activate the burst generators and produce a saccade (Gandhi and Bonadonna, 2005), while the BREM itself reaches completion in parallel. The observation that blinks remove a potent source of inhibition on the oculomotor system has been used to study latent dynamics of movement cancellation (Walton and Gandhi, 2006), characteristics of SC stimulation-evoked saccades under extended disinhibition (Katnani et al., 2012), and the relationship between target selection and motor preparation (Katnani and Gandhi, 2013). In addition to the BREM and saccade components of blink-triggered movements, leakage of premotor activity before saccade onset (but after blink onset) may produce slow eye movements, as seen in Figure 32c (non-zero velocity before t=0 in the blue trace). These observations and those of previous studies, along with the demonstration of linear independence, suggests that the reflex blink may be a simple and elegant tool to probe the dynamics of movement preparation in the saccadic system.

4.0 REMOVAL OF INHIBITION UNCOVERS LATENT SACCADE PREPARATION DYNAMICS

4.1 INTRODUCTION

The ability to interact with the world through movements is a hallmark of the animal kingdom. Movements are usually preceded by a period of planning, when the nervous system makes decisions about what the optimal response to a stimulus is and programs its execution. Such planning behaviour is seen in a wide variety of species, including, insects (Fotowat and Gabbiani, 2007; Card and Dickinson, 2008), fishes (Preuss et al., 2006), frogs (Nakagawa and Nishida, 2012), and mammals (Hanes and Schall, 1996; Churchland et al., 2006a). A fundamental question in sensorimotor neuroscience is how the brain prepares for movements before issuing the command to initiate them.

This question has been studied extensively in the context of gaze control in primates. An influential model of movement preparation and initiation in the gaze control system is the threshold hypothesis (Hanes and Schall, 1996). Inspired by stochastic accumulator models of decision-making (Carpenter and Williams, 1995; Ratcliff and Rouder, 1998), this model posits that saccade initiation is controlled by accumulation to threshold of a motor preparation signal in premotor neurons. Build-up of low frequency activity prior to the saccade command, or its signatures, have been observed in a wide variety of brain regions involved in gaze control, including the frontal eye fields (Hanes and Schall, 1996; Gold and Shadlen, 2000), lateral intraparietal area (Platt and Glimcher, 1999), and superior colliculus (SC) (Dorris et al., 1997), and variability in build-up rate has been shown to be correlated with variability in saccade reaction times. However, it is unclear whether the preparatory activity must rise to a fixed biophysical threshold at the level of individual neurons or a population of neurons in order to initiate a saccade (Hanes and Schall, 1996; Zandbelt et al., 2014), or whether there exists a dynamic equivalent to such a threshold (Lo and Wang,

2006). Indeed, it has recently been shown that a threshold computed under the static assumption can vary across task conditions or epochs, thus challenging the classic interpretation of the threshold hypothesis (Jantz et al., 2013).

Apart from movement preparation, studies with cleverly designed behavioural tasks have attributed the low-frequency build-up of activity to a number of cognitive processes, including target selection (Schall and Hanes, 1993; Horwitz and Newsome, 1999; McPeek and Keller, 2002; Carello and Krauzlis, 2004), attention (Goldberg and Wurtz, 1972; Ignashchenkova et al., 2004; Thompson et al., 2005), decision-making (Newsome et al., 1989; Gold and Shadlen, 2000), working memory (Sommer and Wurtz, 2001; Balan and Ferrera, 2003), and reward prediction (Platt and Glimcher, 1999; Hikosaka et al., 2006). In the presence of such seeming confounds, it is unclear when and how sensory and cognitive processes end and movement preparation begins. One possibility is that neurons in sensorimotor structures represent these signals in parallel, preparing for movements in proportion to the strength of cognitive signals - the premise of the premotor theory of attention (Rizzolatti et al., 1987; Hoffman and Subramaniam, 1995). Alternatively, it is possible that these signals undergo a transformation from representing cognitive processes into movement-related signal at a discrete point in time (Thompson et al., 1996; Horwitz and Newsome, 1999; Juan et al., 2004; Schall et al., 2011), when the activity reaches a movement initiation criterion, thereby acquiring the potential to generate a movement.

How might we test for the presence of such a "motor potential" in low frequency neural activity? The following thought experiment helps illustrate a possible approach. Consider the activity of a neuron rising to an arbitrary threshold. Under normal circumstances, inhibitory gating on the saccadic system is released at a specific time, possibly when activity reaches the purported threshold level, generating a movement (Figure 41a, gray traces). Instead, if inhibition was somehow removed at a prior time through an experimental manipulation (vertical black line in Figure 41a), the occurrence of a movement would indicate that ongoing neural activity possessed motor potential. Furthermore, the dynamics of activity following the manipulation would indicate whether the activity much reach a specific threshold in order to produce the movement (Figures 41b and 41c).

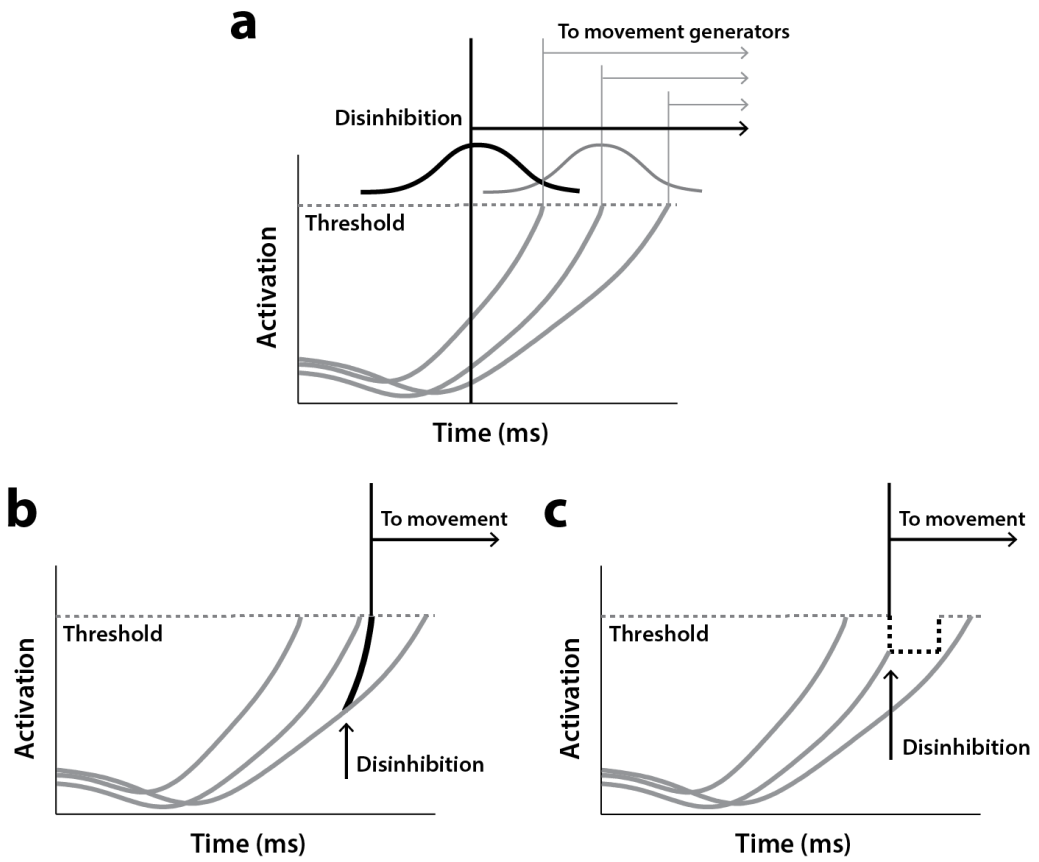

Figure 41. Testing models of movement preparation

a. Schematic representation of the threshold model of movement preparation. Neural activity in premotor networks accumulates at different rates (three gray traces) to a constant threshold (dotted gray line), opening downstream gates and triggering movement generation at different times (gray distribution). Removal of inhibition through an experimental manipulation at earlier times (vertical black line) can trigger movements with reduced reaction times (black distribution) and test specific predictions of the threshold hypothesis. **b.** The fixed threshold model would predict that activity must shoot up (thick black trace) to threshold level following disinhibition (vertical arrow) in order to generate a movement. **c.** Alternatively, it is possible that a movement is triggered off the current level of activity (dotted black reduction from threshold) at the time of disinhibition (vertical arrow).

In this study, we used the trigeminal blink reflex to remove inhibition on the gaze control network during ongoing low-frequency activity. The omnipause neurons (OPNs) in the brainstem, which are tonically active during fixation and turn off during saccades (Cohen and Henn, 1972), also turn off during eye movements associated with blinks (Schultz et al., 2010). Previous work in our lab has shown that removal of this potent source of inhibition on the saccade burst generators with reflex blinks triggers saccades at lower-than-normal latencies (Gandhi and Bonadonna, 2005), an observation that has been used to study dynamics of movement cancellation during saccade countermanding (Walton and Gandhi, 2006), the motor potential of a target selection signal during visual search (Katnani and Gandhi, 2013), and latent sensorimotor processing (Jagadisan and

Gandhi, 2016). Here, we found that it is not necessary for preparatory activity in SC to reach a threshold before a saccade is produced – neural activity just prior to saccades triggered by reflex blinks was lower at both individual neuron and population levels. These movements were also preceded by an acceleration (but not to threshold) of ongoing activity following the manipulation. Moreover, the level of preparatory activity at the time of the blink was correlated with initial dynamics of the evoked movement, suggesting that ongoing sub-threshold activity in SC indeed possesses motor potential.

4.2 METHODS AND DATA ANALYSIS

4.2.1 General and surgical procedures

All experimental and surgical procedures were approved by the Institutional Animal Care and Use Committee at the University of Pittsburgh and were in compliance with the US Public Health Service policy on the humane care and use of laboratory animals. We used two adult rhesus monkeys (*Macaca mulatta*, 1 male and 1 female, ages 8 and 10, respectively) for our experiments. Under isoflurane anesthesia, a craniotomy that allowed access to the SC was performed and a recording chamber was secured to the skull over the craniotomy. In addition, posts for head restraint and scleral search coils to track gaze shifts were implanted. Post-recovery, the animal was trained to perform standard eye movement tasks for a liquid reward.

4.2.2 Visual stimuli and behavior

Visual stimuli were displayed by back-projection onto a hemispherical dome. Stimuli were white squares on a dark grey background, 4x4 pixels in size and subtended approximately 0.5° of visual angle. Eye position was recorded using the scleral search coil technique, sampled at 1 kHz. Stimulus presentation and the animal's behavior were under real-time control with a LabVIEW-based controller interface (Bryant and Gandhi, 2005). After initial training and acclimatization, the monkeys were trained to perform a delayed saccade task. The subject was required to initiate the

trial by looking at a central fixation target. Next, a target appeared in the periphery but the fixation point remained illuminated for a variable 500-1200 ms, and the animal was required to delay saccade onset until the fixation point was extinguished (GO cue). Trials in which fixation was broken before peripheral target onset were removed from further analyses. The animals performed the task correctly on >95% of the trials.

4.2.3 Induction of reflex blinks

On 15-20% of trials, fixation was perturbed by delivering an air puff to the animal's eye to invoke the trigeminal blink reflex. Compressed air was fed through a pressure valve and air flow was monitored with a flow meter. To record blinks, we taped a small Teflon-coated stainless steel coil (similar to the ones used for eye tracking, but smaller in coil diameter) to the top of the eyelid. The air pressure was titrated during each session to evoke a single blink. Trials in which the animal blinked excessively or did not blink were aborted and/or excluded from further analyses. To obtain blink-triggered movements to the peripheral target, air puff delivery was timed to evoke blinks 100-250 ms after the GO cue, during the early phase of the typical saccade reaction time. Blink-triggered movements are typically non-saccadic – they are contaminated by a loopy, blink-related eye movement (BREM) that usually occurs when the animal blinks during fixation. As described in the linear decomposition section below, we extracted the saccadic component of blink-triggered movements by subtracting the BREM from the combined movement. Thus, to obtain BREMs for use in this approach, we also delivered air puffs to evoke blinks during fixation of the central target, 400-100 ms before the onset of the peripheral target. The window constraints for gaze were relaxed for a period of 200-500 ms following delivery of the air puff to ensure that the excursion of the BREM did not lead to an aborted trial.

4.2.4 Electrophysiology

During each recording session, a tungsten microelectrode was lowered into the SC chamber using a hydraulic microdrive. Neural activity was amplified and band-pass filtered between 200 Hz and 5 kHz and fed to a digital oscilloscope for visualization and spike discrimination. A window discriminator was used to threshold and trigger spikes online, and the corresponding spike times

were recorded. The location of the electrode in the SC was confirmed by the presence of visual and movement-related activity as well as the ability to evoke fixed vector saccadic eye movements at low stimulation currents (20-40 μA, 400 Hz, 100 ms). Before beginning data collection for a given neuron, its response field was roughly estimated. During data collection, the saccade target was placed either in the neuron's response field or at the diametrically opposite location (reflected across both axes) in a randomly interleaved manner.

4.2.5 Data analysis and pre-processing

Data were analyzed using a combination of in-house software and Matlab. Eye position signals were smoothed with a phase-neutral filter and differentiated to obtain velocity traces. Normal saccades, blink-related eye movements, and blink-triggered gaze shifts were detected using standard onset and offset velocity criteria (50 deg/s and 30 deg/s, respectively). Onsets and offsets were detected separately for horizontal and vertical components of the movements and the minimum (maximum) of the two values was taken to be the actual onset (offset). Saccade onset times within blink-triggered movements were detected by the linear decomposition method described below.

Raw spike density waveforms were computed for each neuron and each trial by convolving the spike trains with a Gaussian kernel (width = 4 ms). For a given neuron and target location, spike densities were averaged across trials after aligning to saccade onset. We also normalized the trial-averaged spike density of each neuron to enable meaningful averaging across the population. The activity of each neuron was normalized by its peak firing rate during normal saccades.

4.2.6 Linear decomposition of blink-triggered movements

We modeled blink-triggered movements as a linear combination of a BREM and a saccade with an arbitrary delay between the two signals, as described in Chapter 3. The model is described by the following equation:

$$\dot{\boldsymbol{M}}(t) = \dot{\boldsymbol{b}}_j(t) + \dot{\boldsymbol{s}}_k(t - \Delta)$$

where $\boldsymbol{M}(t)$, $\boldsymbol{b}(t)$, and $\boldsymbol{s}(t)$ represent the position signals of the blink-triggered movement, BREM, and normal saccade, respectively, and the overdot denotes the derivative of those signals, representing velocity. We used a least-squares fit of the simulated blink-triggered movement generated by the model above to recorded blink-triggered movements to optimize the value of the delay, Δ. This value was taken as the time of saccade onset, relative to blink-triggered movement onset, and was used in all subsequent analyses. Saccade onset was typically delayed with respect to BREM onset in blink-triggered movements.

4.2.7 Surrogate data analysis

For the accumulation rate analysis in Figure 44, we created a surrogate dataset of control trials with blink times randomly sampled from the distribution of blink occurrences in perturbation trials for that session and assigned to individual control trials. For each neuron, we created 1000 such pseudo-trials by resampling from and reassigning to control trials. We only used blink trials for which saccade onset was delayed with respect to blink onset by atleast 30 ms. This restriction reduced the number of neurons for which sufficient data (at least 10 trials) was available to 32, and provided us with a 20 ms window for accumulation rate analysis following the blink before a putative threshold time of 20 ms before saccade onset. We then fit the accumulation rates 20 ms before and after the blink with piecewise-linear functions and compared the change in accumulation rates before and after the pseudo-blink in control trials and the blink in perturbation trials.

4.3 RESULTS

In order to explicitly test the threshold hypothesis within a single behavioural paradigm (and therefore matched states of the premotor networks involved in gaze control), we induced reflex blinks after the go cue in monkeys performing the delayed saccade task. As mentioned in the Introduction, a reflex blink is a suitable perturbation because it removes inhibition on the saccadic system by turning off the OPNs and triggers gaze shifts at lower-than-normal latencies (Gandhi

and Bonadonna, 2005). We have previously described a method to extract saccades from blink-triggered eye movements by linear decomposition (Chapter 3). We computed onset times of saccades embedded in blink-triggered movements using this method.

4.3.1 The blink perturbation triggers reduced latency saccades

We first verified that this combination of reflex blinks and the linear decomposition method of detecting saccade onset produced low-latency saccades in accordance with previous observations. Figure 42 shows saccade reaction time (from GO cue) as a function of the time of blink in perturbation trials (light and dark blue circles). To visually compare reaction times on blink trials with those in control trials, it was necessary to include the distribution of control reaction times in this figure. To do this, we created a surrogate dataset similar to the one used for the analysis in Figure 44 by randomly assigning blink times to control trials (see Methods), and plotted them on the same axes as blink trials in Figure 42 (black circles).

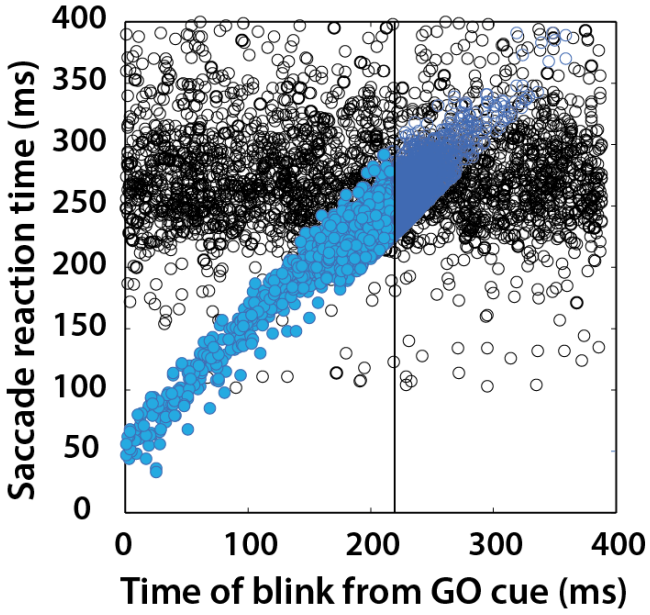

Figure 42. Time course of blink-triggered saccades

Saccade reaction time as a function of blink onset time. Open and filled blue circles are individual blink trials, and black circles are from control trials with randomly assigned blink times for comparison. The rest of the analysis in this study focuses on the early preparatory period and therefore only includes trials (filled light blue circles) where the blink occurred <220 ms (vertical line) after the GO cue.

Reaction times in perturbation trials were correlated with time of blink, and were significantly lower than control reaction times, consistent with previous observations (Gandhi and Bonadonna, 2005). Since our goal was to study movement preparation dynamics following disinhibition of the saccadic system by the blink, as a way to gain insights into models of movement initiation, we wanted to ensure that the activity in the network was sufficiently far away from a mature motor command. Therefore, the rest of the analyses in this manuscript are focused on the trials when the blink was evoked relatively early compared to typical reaction times in control trials (< 220 ms, filled light blue circles to the left of vertical line in Figure 42).

4.3.2 Blink-triggered saccades are evoked at lower thresholds compared to normal saccades

Next, we asked whether it is necessary for activity in SC intermediate layer neurons to reach a fixed threshold in order to generate a movement. Previous studies have estimated the threshold for individual neurons by assuming a specific time at which the threshold could be reached before saccade onset or by computing the time, backwards from saccade onset, at which premotor activity starts becoming correlated with reaction time (Hanes and Schall, 1996). Given the heterogeneity of the activity profiles of premotor neurons, we think these approaches are too restrictive and not suitable for an unbiased estimate of the threshold. Instead, we took a non-parametric approach adopted previously (Jantz et al., 2013) and scanned through possible times at which threshold might be reached prior to saccade onset. Figure 43a shows a snippet of the average normalized population activity aligned on saccade onset for control (black traces) and blink (blue traces) trials. For each neuron in this population (n = 48), we computed the average activity in 10 ms bins slid in 10 ms increments from 50 ms before to 20 ms after saccade onset (coloured windows at the bottom of Figure 43a). If activity on control trials reaches the purported threshold at any one of these times before saccade onset, a comparison with activity in blink trials at that time should reveal the existence, or lack thereof, of a fixed threshold. Figure 43b shows the activity in these bins for control trials plotted against blink trials, coloured according to the bins in Figure 43a. Note that the majority of points for early time bins lie below the unity line. Activity on blink trials was significantly lower compared to control activity from 50 ms before to 0 ms before saccade onset (square points, Wilcoxon signed-rank test, $p < 0.0001$). The systematic trend in the linear fits (solid

lines) to these points suggests that the activity on blink trials gradually approaches that on control trials; however, the earliest time at which control activity was not different from activity in blink trials was 10 ms after saccade onset (circles) – too late to be considered activity pertaining to a movement initiation threshold. Thus, activity at the population level need not reach a threshold level in order to produce a movement.

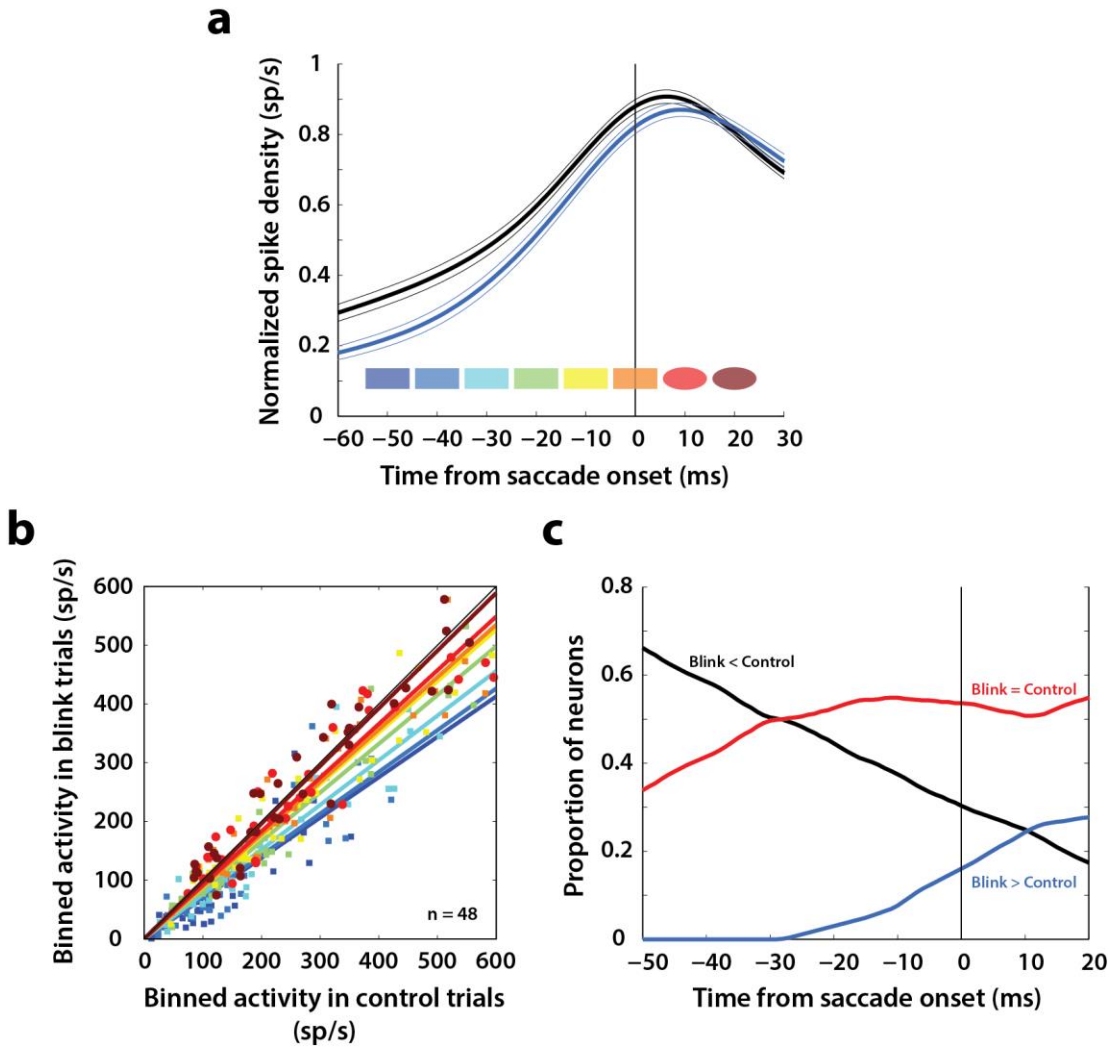

Figure 43. Analysis of putative threshold

a. Average normalized population activity (thick traces) aligned on saccade onset for control (black trace) and blink (blue trace) trials. The thin lines represent s.e.m. The coloured swatches at the bottom show the time windows used for the analysis presented in **b**; their shapes represent the presence (rectangle) or absence (ellipse) of a significant difference between control and blink rates in that time window. **b.** Scatter plot of the activities of individual neurons in the time windows illustrated in **a** for control versus blink trials. The diagonal (thin black line) is the unity line. Square points indicate that the activity in control trials was higher in that window compared to blink trials, and circles indicate that there was no significant difference between the two conditions. Coloured lines indicate linear fits to the scatter at the corresponding time window. **c.** Proportion of neurons exhibiting the labelled differences between the two conditions as a function of time.

Nevertheless, we wanted to know if there exist individual neurons in the population that might obey the threshold hypothesis. For each neuron, we calculated whether activity on blink trials was higher, lower, or not significantly different from activity in control trials, at each time point from -50 before to 20 ms after saccade onset (Wilcoxon rank-sum test, $p<0.001$). The three traces in Figure 43c represent the proportion of neurons that showed each of those three characteristics as a function of time. As late as 10 ms before saccade onset, roughly 40% of the neurons had lower activity on blink trials compared to control trials (black trace), inconsistent with the idea of a fixed threshold. Slightly over half the neurons did not exhibit significant differences in activity on blink and control trials at that time point (red trace); however, this observation is insufficient to conclude that the activities in the two conditions were identical, or that it must reach a threshold. Of course, it is possible that some of these neurons belong to a class for which fixed thresholds have been observed in previous studies. Together, these results suggest that it is not necessary for premotor activity in SC intermediate layers to reach a threshold level at the individual neuron or population level in order to produce a movement.

4.3.3 Rate of accumulation of SC activity accelerates following disinhibition by the blink

We mentioned in the Introduction that in order for neurons to reach a fixed threshold following removal of inhibition and generate reduced latency saccades, the rate of accumulation of the preparatory signal must increase (Figure 41c). As we saw above, SC neurons do not necessarily have a fixed threshold. It is nonetheless possible that the dynamics of preparatory activity is modulated when the network is perturbed by a blink, possibly through feedback interactions between the brainstem circuitry and SC. Since we wanted to test for a change in dynamics before the actual saccade started, we restricted our analysis to the subset of trials in which saccade onset occurred at least 30 ms after blink onset. This restriction reduced our population to 32 neurons. For each neuron, we estimated the rate of accumulation of activity in 20 ms windows before and after blink onset with piecewise linear fits (Figure 44a, dotted blue trace and solid lines).

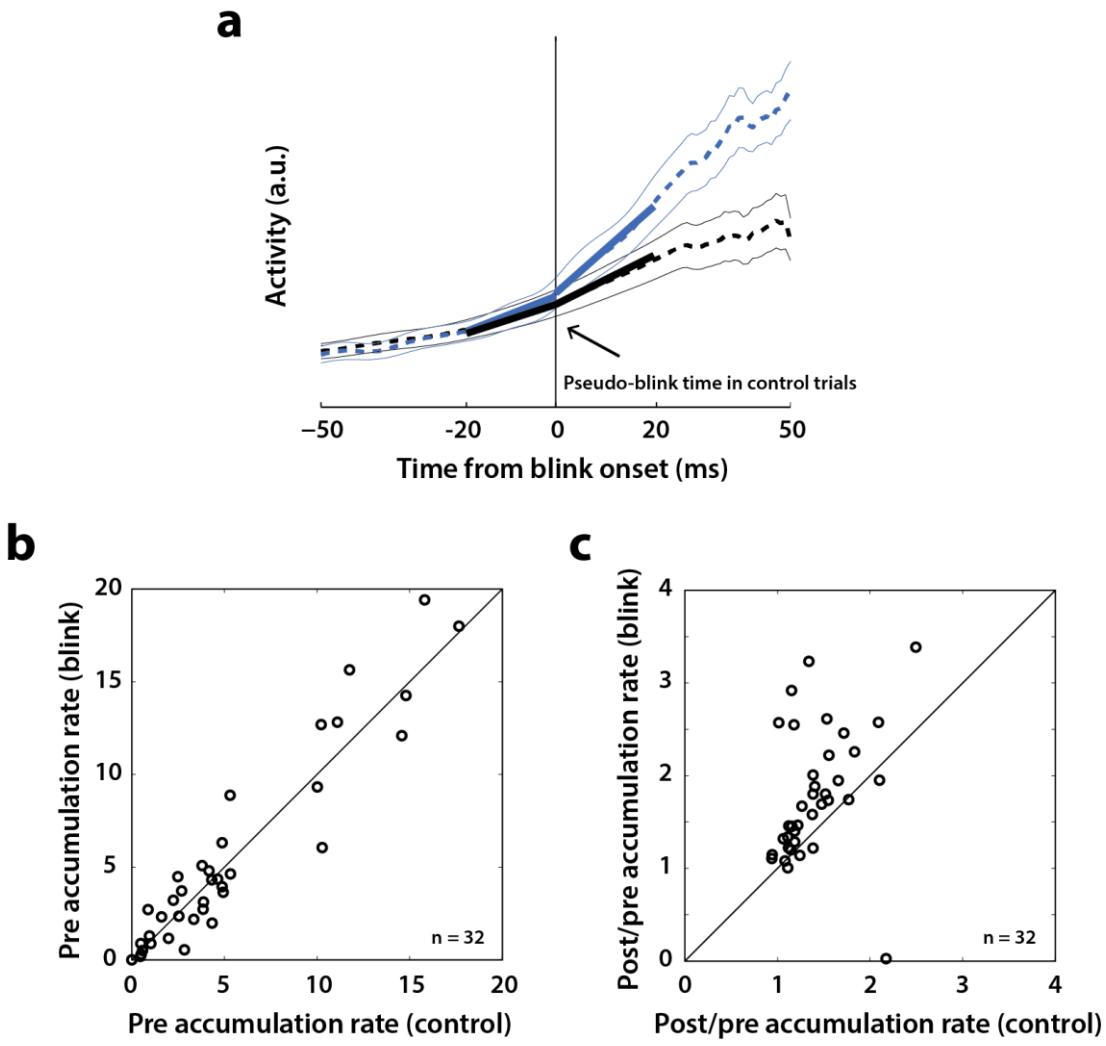

Figure 44. Analysis of accumulation rate change following perturbation

a. Schematic illustrating the computation of accumulation rates before and after the blink. The snippet shows the average population activity centered on blink onset for blink trials (blue trace) and pseudo-blink onset from the surrogate dataset for control trials (black trace). The thick lines represent linear fits to the activity 20 ms before and after the blink time. **b.** Scatter plot of pre-blink accumulation rates of individual neurons for control versus blink trials. The unity line is on the diagonal. **c.** Scatter plot of the ratio of post-blink to pre-blink rates for control versus blink trials.

It is important to note that while the evolution of premotor activity is commonly modelled as a linear process, the actual dynamics of accumulation may be non-linear, causing spurious changes in linear estimates of accumulation rate over time. To account for this, we created a surrogate dataset of control trials for each neuron, with blink times randomly assigned from the actual distribution of blink times for that session. We also performed the same procedure on the control dataset (dotted black trace and solid lines in Figure 44a). Changes in accumulation rate on control trials following the pseudo-blink should reflect the natural evolution of activity at typical

blink times and provide a baseline for comparing any changes observed in blink trials. We first verified that the accumulation rates on control and blink trials were matched before the blink. Figure 44b shows a scatter plot of pre-blink linear slopes for 32 neurons on control and blink trials; rates were not different between the two conditions (Wilcoxon signed-rank test, $p > 0.5$). Next, we tested for a change in accumulation rates following the blink by calculating the ratio of the post- and pre- blink linear fits (Figure 44c) in each condition. For all neurons, this ratio was greater than 1 even on control trials, highlighting the natural non-linear dynamics mentioned above. The change in accumulation rate was significantly higher following the actual blink on blink trials (Wilcoxon signed-rank test, $p < 0.0001$). Thus, removal of inhibition seems to cause an acceleration in the dynamics of ongoing activity in the lead up to a saccade.

4.3.4 SC preparatory activity preceding the saccade-related burst possesses motor potential

Finally, we wanted to know whether ongoing preparatory activity possessed any motor potential before its maturation into a motor command. We hypothesized that, if the low-frequency activity had motor potential, removing inhibitory gating on the system will result in a slow eye movement proportional to the level of activity, before accelerating into a saccade in accordance with acceleration of the neural activity. Recall that the underlying saccade is recovered from the blink-triggered movement by subtraction of a BREM at an optimal delay (Methods). This typically results in a residual velocity signal before saccade onset (Figure 45a, blue trace). We considered the possibility that the residual velocity was the result of an actual drift of the eyes, possibly as a manifestation of the motor potential of low-frequency activity. On the other hand, if the residual velocity were unrelated to motor potential, there should be no correlation with neural activity.

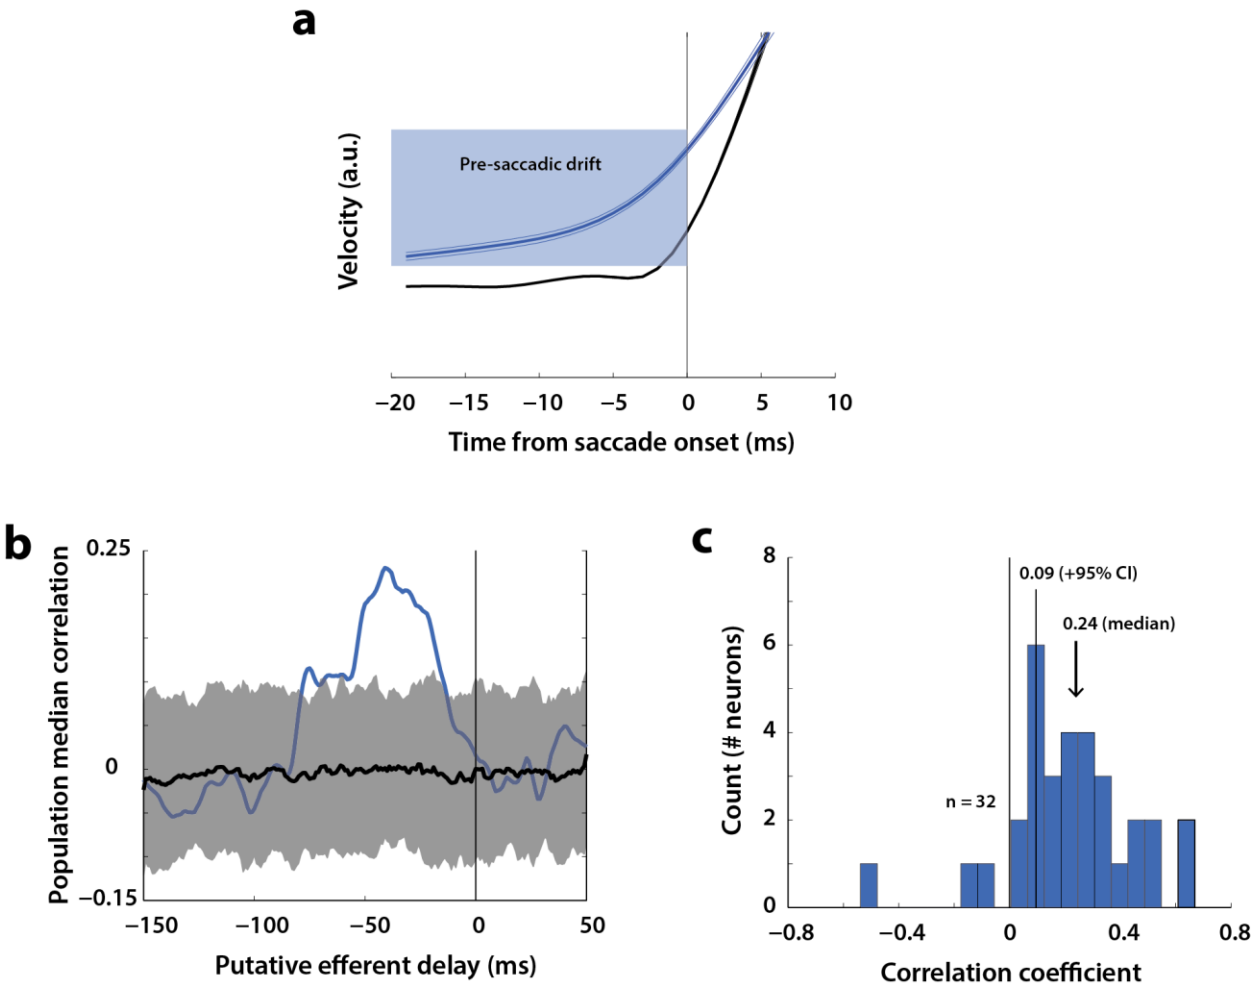

Figure 45. Correlation between neural activity and pre-saccadic velocity

a. Schematic illustrating the pre-saccadic drift and the time window used for calculating velocity. The snippet shows the across-session average of radial velocity before saccade onset for control trials (black trace) and blink trials (blue trace). **b.** Correlation between velocity and neural activity as a function of lag of the activity window with respect to velocity. Negative values of the lag ("putative efferent delay") represent activity windows shifted earlier in time compared to the velocity window. Correlations were computed for each neuron between single trial correlates. The blue trace is the median correlation across the population, the black trace is the bootstrapped (trial-shuffled correlation) median and the gray region is the +/- 95% confidence interval for the bootstrap. **c.** Distribution of correlation coefficients at the optimal efferent delay of 40 ms. The vertical line at 0.09 is the +95% CI value from **b**.

To test this, for each neuron, we calculated the correlation across trials between pre-saccadic velocity in a 20 ms window (shaded blue region) and neural activity in a corresponding 20 ms window. To account for efferent delays between SC activity and the extraocular muscles, we adopted a non-parametric approach and performed this analysis on a range of possible time lags between neural activity and velocity. Figure 45b shows the median correlation of the population of 32 neurons as a function of putative efferent delay (blue trace; the black traced and shaded region represent the trial-bootstrapped median and 95% CI, respectively). Pre-saccadic

velocity was significantly correlated with activity 55 ms to 20 ms prior, revealing the substantial presence of a motor potential in the low-frequency activity of SC neurons. The occurrence of significant correlations for a range of efferent delays suggests the possibility that activity may be integrated over time to produce the eye movement. The distribution of correlations for the optimal putative delay (40 ms) is shown in Figure 45c. Note that the majority of neurons (24/32, 75%) lie above the 95% confidence interval obtained from bootstrap analysis, suggesting that their activities possess latent motor potential. Thus, ongoing activity in SC neurons, well before typical saccade times, encodes an evolving motor command, and the level of this activity represents the immediacy of the potential to initiate a saccade.

4.4 DISCUSSION

In this study, we sought to uncover the dynamics of movement preparation and, specifically, test an influential model of saccade initiation – the threshold hypothesis (Hanes and Schall, 1996). By disinhibiting the saccadic system much earlier than its natural time course by inducing a reflex blink, we showed that low-frequency preparatory activity in the intermediate layers of SC does not have to reach a threshold at the single neuron or population level in order to initiate a saccade. We found that the rate of accumulation of activity increases following disinhibition before the saccade is initiated. Furthermore, we found that the low-level activity in premotor neurons has a latent "motor potential", which is uncovered when downstream gating is turned off.

4.4.1 Implications for threshold-based accumulator models

Other recent studies have argued that the fixed threshold hypothesis does not hold true for most neurons in SC and FEF (Jantz et al., 2013), finding that the effective threshold varies based on the task being performed by the subject. However, comparison of thresholds across tasks is subject to the confound that the network may be in a different overall state, thereby modulating the threshold. For instance, the presence or absence of the fixation spot at the time of movement initiation, or the presence of other visual stimuli or distractors, may affect how downstream neurons receiving

premotor activity from the whole network decode the level of activity, thus influencing the effective threshold. In our study, thresholds are compared between interleaved blink and control trials in the same behavioral paradigm, reducing the effect of network-level confounds. Admittedly, the reflex blink is designed to trigger saccades by turning off the OPNs (Schultz et al., 2010), which provide strong inhibitory gating on these movements downstream of the SC (Cohen and Henn, 1972). It is possible that the effective reduction in threshold that we observe is due in part to reduced OPN activity – premotor activity has to overcome lower inhibition and therefore triggers movements at a lower level. It is also important to note that while there is some evidence that premotor activity in SC is attenuated when saccades are perturbed by a reflex blink (Goossens and Van Opstal, 2000b), we did not observe suppression during movements that were triggered by the blink, as seen in the firing rate profile in Figure 44a. Nevertheless, the presence of any attenuation would only strengthen the result, since the occurrence of saccades despite attenuated SC activity goes against the notion of a rigid threshold.

4.4.2 Alternative models of movement initiation

It has also been suggested that the threshold may not be rigid at the single neuron level but may instead be implemented by pooling the activity of the premotor neuron population (Zandbelt et al., 2014). However, population-level thresholding does not seem like a necessary condition either, based on the results here (Figure 43a). If activity in premotor structures does not reach a specific threshold in order to initiate a movement, how do burst generators in the brainstem know when to initiate a movement? One possibility is that the preparatory signal must cross an inferred decision bound (as opposed to a biophysical bound) of the decoder, which, in general, may be a non-linear function of population activity. Indeed, studies of movement preparation in the skeletomotor system have shown that neural activity reaches an optimal subspace before undergoing dynamics that produce a limb movement (Afshar et al., 2011; Churchland et al., 2012). A related idea suggests that activity pertaining to movement preparation evolves in a region of population space that is orthogonal to the optimal subspace, and this dissociation confers neurons the ability to prepare the movement by incorporating perceptual and cognitive information without risking a premature movement (Kaufman et al., 2014). However, just as with the threshold hypothesis, it is unclear whether these correlative observations are epiphenomena or causally linked to the process

of initiating movements. For example, what aspects of population activity in the null space prevent it from releasing downstream inhibitory gating while permitting its flow to produce movements when in the optimal subspace? More studies that delink evolving population activity from physiological gating are needed to clearly delineate these mechanisms.

4.4.3 Parallel implementation of the sensory-to-motor transformation

An influential idea in systems neuroscience is the premotor theory of attention, which posits that spatial attention is manifested by the same neural circuitry that produces movements (Rizzolatti et al., 1987). Low-frequency preparatory activity in SC has been attributed to a number of cognitive processes including spatial attention (Goldberg and Wurtz, 1972; Ignashchenkova et al., 2004) and target selection (Horwitz and Newsome, 1999; McPeek and Keller, 2002), but the question of whether this signal exclusively represents those processes or whether it represents the intention to shift gaze in parallel has been debated but unresolved. Efforts to delineate spatial attention and movement intention by means of causal manipulations have produced a mixed bag of results, with some studies supporting disjoint processing (Juan et al., 2004) and others supporting parallel processing (Moore and Armstrong, 2003; Katnani and Gandhi, 2013). The discovery of a latent motor potential in the preparatory activity of SC neurons contributes to this debate by suggesting that while the low-frequency build-up may not trigger movements under normal conditions, movement intention and motor programming signals are also encoded by these neurons in parallel. Moreover, unlike manipulations such as microstimulation or pharmacological inactivation that introduce extrinsic signals that may corrupt the natural processing of this activity, reflex blinks are non-invasive and are likely to provide a more veridical readout of ongoing processes.

It is worthwhile to end on a note of caution. The results in this study are based on experiments performed in one node, SC, in a distributed network of brain regions involved in gaze control. Traditional knowledge imposes a hierarchy on the sensorimotor transformations that need to occur before a gaze shift is generated (Wurtz et al., 2001). It is possible that sensorimotor neurons in SC, and to some extent, FEF, which project directly to the brainstem burst generators (Segraves, 1992; Rodgers et al., 2006), are more likely to exhibit signatures of a motor potential in preparatory activity compared to regions higher in the cascade. Neurons in other regions may

still need to signal the initiation of a movement by reaching a threshold, optimal subspace, or other similar decision bound. Future studies that take causal approaches to perturbing intrinsic population dynamics in these regions are essential in order to gauge whether the findings in this study are specific to one structure or a more generalizable property of sensorimotor systems.

5.0 POPULATION TEMPORAL STRUCTURE SUPPLEMENTS THE RATE CODE DURING SENSORIMOTOR TRANSFORMATIONS

5.1 ABSTRACT

In order to successfully interact with the environment the brain must funnel down the sensory inputs it receives to specific movements at specific times. Such sensory-to-motor transformations are critically mediated by neurons in premotor brain networks whose evolving activities represent sensory, cognitive, and movement-related information (Gallese et al., 1996; Wurtz et al., 2001; Buneo et al., 2002). However, this multiplexing poses a conundrum – how does a decoder know precisely when to initiate a movement if such neurons are also active at times other than when a movement occurs (e.g., in response to sensory stimulation)? Extant models of movement generation that rely on firing rate, including rise-to-threshold (Hanes and Schall, 1996), inhibitory gating (Keller, 1974), and dynamical switches at the population level (Churchland et al., 2012; Kaufman et al., 2014), leave certain explanatory gaps unfilled. Here, we propose a novel hypothesis: movement is triggered not only by an increase in firing rate, but critically by a reliable temporal pattern in the population response. We show that in brain regions involved in orienting eye movements - the superior colliculus (SC) and the frontal eye fields (FEF) - the temporal dynamics between neurons during visually-driven activity is different from that during pre-movement activation. Specifically, using a measure that captures the fidelity of the population code - here called temporal stability - we show that the temporal structure fluctuates in the visual response but remains stable during the pre-movement response, thus distinguishing incoming sensory input from motor output. This is an important attribute because SC and FEF "visuomovement" neurons project directly to the brainstem (Segraves, 1992; Rodgers et al., 2006) which houses the controller for gaze shifts, and any increase in the incoming drive is poised to trigger a (potentially undesirable) gaze shift. We also demonstrate that a simple firing rate-based

network with synaptic facilitation can distinguish between stable and fluctuating population codes, suggesting a putative mechanism for interpreting temporal structure. These findings offer an alternative perspective on the relationship between spatial attention and motor preparation (Rizzolatti et al., 1987; Kustov and Robinson, 1996) by situating the correlates of movement initiation in the temporal features of activity in shared neural substrates. They also suggest a need to look beyond the instantaneous rate code and consider the effects of short-term population history on neuronal communication and its effects on behaviour.

5.2 INTRODUCTION

Across brain regions involved in the control of gaze, including the SC in the midbrain and the FEF in the neocortex, so-called visuomovement neurons burst a volley of spikes both in response to a visual stimulus and for generating a gaze shift or saccade to the spatial location of the stimulus (Wurtz et al., 2001). The dual nature of visuomovement neurons is best illustrated by examining their activity in the delayed response paradigm (Figure 51A). In this task, the subject is required to withhold a saccade in response to the presentation of a stimulus in the visual periphery until after the disappearance of a central fixation cue. Figure 51C shows the average activity aligned on target (left) and saccade (right) onsets for 57 SC and 22 FEF neurons in two monkeys (*Macaca mulatta*). The visual response exhibits a high frequency burst, reaching levels observed 30-40 ms prior to saccade onset, yet the former rise does not trigger a movement while the latter does. In fact, it is not uncommon to find neurons whose peak visual response is greater than peak pre-movement activation (Jantz et al., 2013), an observation that casts doubt on thresholding (Hanes and Schall, 1996) as a singular mechanism. Given that such neurons project directly to the brainstem saccade burst generator that initiates and guides saccadic gaze shifts (Segraves, 1992; Rodgers et al., 2006) (Figure 51B), we asked how downstream structures are able to differentiate between the two bursts.

Previous studies have examined the across-trial reliability of neural responses as a correlate of behavioural states (Churchland et al., 2010; Steinmetz and Moore, 2010; Chang et al., 2012), including movement preparation (Churchland et al., 2006b). We computed the Fano factor, a ratio

of the spike count variance to the mean across trials. In both SC and FEF, the Fano factor decreased prior to saccade onset (Figure 51D), but showed a mixed pattern during the visual burst following stimulus onset (decrease in FEF and a transient increase in SC). Thus, spike count reliability alone is insufficient to explain the lack of a saccade during the high-frequency visual burst. More importantly, since reliability is necessarily computed across trials, the brain cannot use this information to make movement decisions on any given trial. Below, we present a novel hypothesis based on population temporal dynamics to explain the dissociation between the visual burst and behavior.

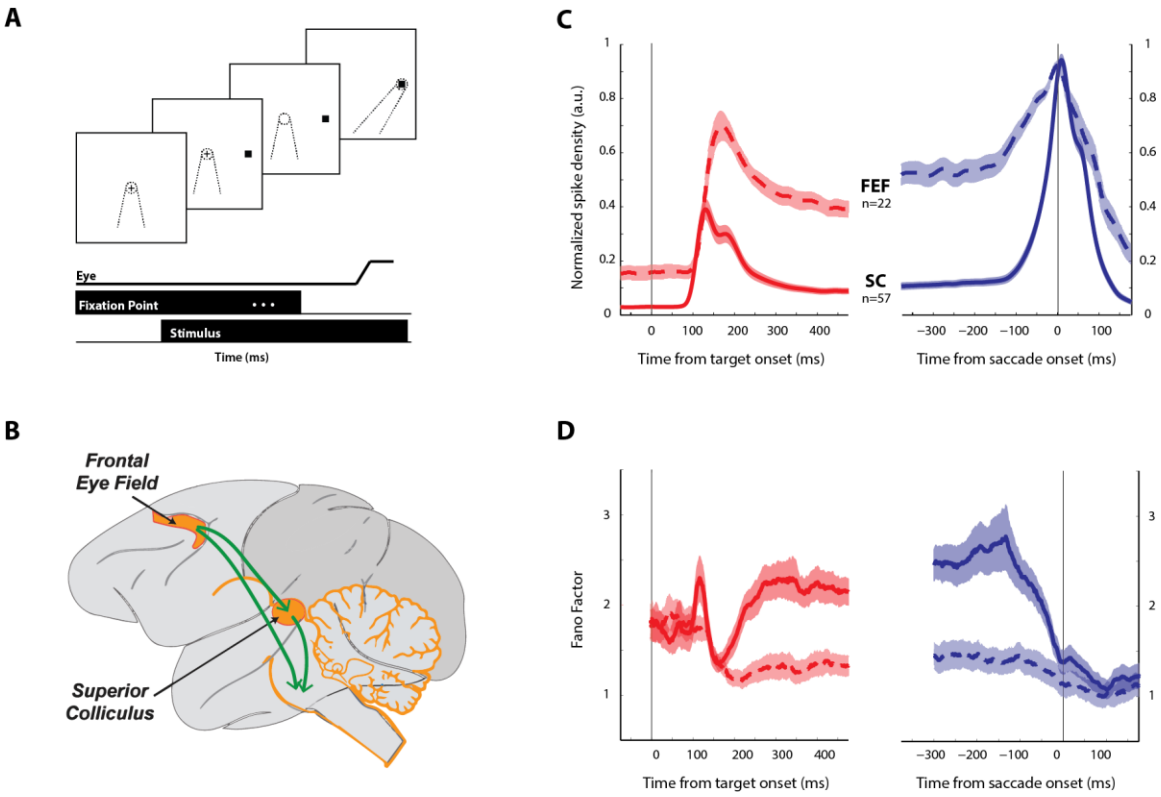

Figure 51. Sensorimotor transformations are mediated by neurons that multiplex sensory and motor information

A. Sequence of events in the delayed saccade task. The top panel shows a typical display sequence and the bottom panel shows the timeline. The fixation target is shown as a plus symbol for illustration purposes. Dotted lines depict line of sight of the animal. **B**. Schematic of the monkey gaze control network showing projections from the FEF and SC to the saccade burst generators in the pons. **C**. Normalized mean population activity of visuomovement neurons recorded in the SC (solid) and FEF (dashed) during the delayed saccade task. Since SC neurons have firing rates ranging from less than 100 spikes/s to 800 spikes/s, we normalized the spike density of each neuron by its peak magnitude to enable meaningful averaging across the population, mainly for display purposes. The left panel shows data aligned to the onset of the target. The right panel shows data aligned to saccade onset. This colour scheme (red and blue for the two alignments respectively) will be followed throughout the article. Shaded regions indicate s.e.m. **D**. Fano factor averaged across the population of SC and FEF neurons. Line types/colour scheme same as in C.

5.3 METHODS

5.3.1 General and surgical procedures

All experimental and surgical procedures were approved by the Institutional Animal Care and Use Committee at the University of Pittsburgh and were in compliance with the US Public Health Service policy on the humane care and use of laboratory animals. We used three adult rhesus monkeys (*Macaca mulatta*, 2 male, ages 8 and 6, and 1 female, age 10) for our experiments. Both SC and FEF were recorded in monkey BB whereas only one of either SC or FEF was recorded in monkeys WM and BL, respectively. Under isoflourane anaesthesia, recording chambers were secured to the skull over craniotomies that allowed access to the SC and FEF. In addition, posts for head restraint and scleral search coils to track gaze shifts were implanted. Post-recovery, the animal was trained to perform standard eye movement tasks for a liquid reward.

5.3.2 Visual stimuli and behavior

Visual stimuli were displayed either by back-projection onto a hemispherical dome (monkeys BB and WM), or on a LED-backlit flat screen television (monkey BL). Stimuli were white squares on a dark grey background, 4x4 pixels in size and subtended approximately 0.5° of visual angle. Eye position was recorded using the scleral search coil technique, sampled at 1 kHz. Stimulus presentation and the animal's behaviour were under real-time control with a Labview-based controller interface (Bryant and Gandhi, 2005). All monkeys were trained to perform standard oculomotor tasks. In the delayed saccade task, the monkey was required to initiate the trial by acquiring fixation on a central fixation target. Next, a target appeared in the periphery but the fixation point remained illuminated for a variable 500-1200 ms, and the animal was required to delay saccade onset until the fixation point was extinguished (GO cue). In the gap task, initial fixation on a central target was followed by a gap period (200 ms) during which the fixation point disappeared, while the animal was required to maintain fixation at the now empty location. This was followed by the appearance of a saccade target in the periphery which was also the GO cue for the animal to make a saccade. All animals performed these tasks with >95% accuracy. Incorrectly performed trials were removed from further analyses. The delayed and gap saccade

tasks were occasionally interleaved with visual search paradigms used for a different study. The gap task, when performed in isolation with predictable spatial and temporal features, is often characterized by a bimodal distribution of saccade reaction times (SRTs) with a proportion of saccades occurring at significantly short latencies (Fischer and Boch, 1983). Since we interleaved several behavioral tasks, the animal was prevented from constantly being in "gap task mode", and the bimodality of the SRT distribution was not as noticeable. Hence, we divided the gap task trials into two sets based on SRT and labelled the set with shorter SRTs "short latency" saccades (100-150 ms) and the other set as "regular" saccades (170-250 ms). The temporal features of the visual and motor bursts in SC neurons during the short latency trials resembled the profiles observed during express saccades (Edelman and Keller, 1996).

5.3.3 Electrophysiology

During each recording session, a tungsten microelectrode was lowered into the FEF or SC chamber using a hydraulic microdrive. Neural activity was amplified and band-pass filtered between 200 Hz and 5 kHz and fed to a digital oscilloscope for visualization and spike discrimination. A window discriminator was used to threshold and trigger spikes online, and the corresponding spike times were recorded. Both SC and FEF were confirmed by the presence of visual and movement-related activity as well as the ability to evoke fixed vector saccadic eye movements at low stimulation currents (20-40 μA, 400 Hz, 100 msec). Before beginning data collection for a given neuron, its response field was roughly estimated. In most cases, during collection of data, the saccade target was placed either in the neuron's response field or at the diametrically opposite location in a randomly interleaved manner. In addition, the simulation-evoked saccades were recorded to identify the response field centers (or "hotspots") for the cells recorded during that session. For recordings in rostral SC, stimuli were presented at one of four locations at an eccentricity sufficient to induce a reduction in activity during the large amplitude saccade.

5.4 DATA ANALYSIS

5.4.1 Preliminary analyses

Data were analyzed using a combination of in-house software and Matlab. Eye position signals were smoothed with a phase-neutral filter and differentiated to obtain velocity traces. Saccades were detected using standard velocity criteria. The animal was considered to be maintaining fixation if the gaze remained within a 2-3° window around the fixation target. We also detected any microsaccades that occurred during the delay period in each trial by using a velocity criterion based on the noise in the velocity signal for that trial. Only one of the two monkeys (WM) in whom we recorded neural activity in the rostral SC made sufficient number of microsaccades to permit further analysis.

Raw spike density waveforms were computed for each neuron and each trial by convolving the spike trains with a Gaussian kernel (width = 4ms; in some instances, we used 10 ms for display purposes only). The spike densities were averaged across condition-matched trials (same target location) following alignment with target or saccade onset. We refer to this averaged spike density function for the i^{th} neuron as $R_i(t)$ below. Since SC neurons have firing rates ranging from less than 100 spikes/s to 800 spikes/s (Figure 52A), we also normalized the average spike density of each neuron by its peak magnitude to enable meaningful averaging across the population, mainly for display purposes (e.g., Figure 51C). Neurons were classified as visuomovement neurons if the spike density was significantly elevated above baseline during the visual epoch (50-200 ms following target onset) and during the premotor epoch (50 ms before and after saccade onset). This classification yielded our data set of 57 neurons in caudal SC and 22 neurons in FEF. In addition, rostral SC neurons were defined as fixation-related if the activity during the premotor epoch of large saccades was significantly reduced below baseline (29 neurons). A subset of these neurons also elevated their discharge around the onset of microsaccades (7 neurons). To minimize the effect of noise in the spike density waveforms due to insufficient number of trials in our analysis, we used only neurons which had at least 10 trials for a given condition. This was not a factor in most conditions (we typically had 50-100 trials) except in the case of the gap task. Applying this

minimum number of trials criterion resulted in 9 neurons for the short latency condition and 41 neurons for the regular latency condition.

We also computed the Fano factor - the spike count variance across trials divided by the mean spike count - for each neuron in the delayed saccade task. Spikes were counted in 100 ms windows sliding in increments of 1 ms. The Fano factor trace was then shifted forward in time by half the window size (50 ms). The rationale behind this correction is that the beginning of any change in the spike count statistics at a given time t is going to be captured in a window centered at t-50 ms. The Fano factor was then averaged across neurons to obtain a population trace (Figure 51D).

5.4.2 Inferring population dynamics from single-unit recordings

Our analyses and results build on the ability to infer features of population dynamics from single-unit activity profiles recorded sequentially over multiple sessions. In order to do this, we constructed a simulacrum of the population activity in SC and FEF by combining all of our single-unit recordings. We define $\mathbf{R}(t)$ as a population activity vector:

$$\mathbf{R}(t) = \begin{bmatrix} R_1(t) \\ R_2(t) \\ \dots \\ R_n(t) \end{bmatrix}$$

where $\mathbf{R}(t)$ represents the instantaneous activity at time t as a point in an *n*-dimensional space, where *n* is the number of neurons in the population. The curve connecting successive points over time is the neural trajectory that describes the evolution of population activity. Note that each neuron's firing rate waveform $R_i(t)$ is the average across many matched trials (identical stimulus/response conditions). Thus the neural trajectory is the *expected trajectory* of population activity. Since these neurons have a fairly broad RF, many neurons constitute the active population for any given visual stimulus or saccade (Sparks et al., 1976). Our neurons were sampled roughly (but not exactly) around the hotspot of the active population for a given session. Therefore, the pooled data from individual sessions can be thought of as an approximation of the population mound active for an arbitrary location/RF in the visual field on any given trial. Many recent studies have reconstructed such "pseudo-populations" from sequentially recorded neurons (Mante et al.,

2013; Rigotti et al., 2013; Stokes et al., 2013) and found comparable properties from simultaneous and serial recordings (Kaufman et al., 2014). Indeed, this is also expected of our pseudo-population under the assumption of isotropy – that each neuron's contribution to its respective local active population is similar regardless of the locus of the population. To better demonstrate this, we estimated the location of a given neuron in the active pseudo-population as follows (we use SC for illustrative purposes because of its convenient topography). We used the point image of the target location on the SC map as a representation as the center of the active population for that session, and used the stimulation-evoked saccade vector to identify the location of that neuron on the SC. We referenced the point image of each target location to a single location to create an active pseudo-population and translated the neuron locations relative to this population center (Figure A1). All mathematical equations for transforming between visual and SC tissue coordinates have been defined previously (Ottes et al., 1986). Inclusion of stimuli/saccades in the anti-preferred RF of the neurons allowed us to also estimate a pseudo-population of neurons in the ipsilateral SC. To complete the representation of activity across the SC topography, we also recorded from and included in our analyses neurons in the rostral portion of SC, which are active during fixation, burst during microsaccades, and are suppressed during large saccades. Moreover, to account for correlations between neurons on individual trials (noise correlation), we show that the temporal stability profile derived from the across-trial mean population dynamics (see next section) is a theoretical upper bound on the stability expected on single trials (Supplementary Material and Figure A1). The findings in this letter are therefore independent of our methodological assumptions and are robust in the face of correlations within the population. It is also important to emphasize that these results are also independent of any across-trial jitter in the phase of the display system's refresh cycle since we use the trial-averaged neural activity for all our analyses.

5.4.3 Temporal stability analyses

To assess temporal stability, we first normalized the population trajectory $\mathbf{R}(t)$ by its Euclidean norm ($\|\mathbf{R}(t)\|$), equivalently its magnitude, at each time point to yield $\widehat{\mathbf{R}}(t)$:

$$\widehat{\mathbf{R}}(t) = \frac{\mathbf{R}(t)}{\|\mathbf{R}(t)\|}$$

The normalized trajectory can be visualized as a unity length population vector that points in an n-dimensional direction at each instant in time. That is, while $\mathbf{R}(t)$ is free to traverse the n-dimensional activity space, $\widehat{\mathbf{R}}(t)$ is constrained to the surface of an n-dimensional hypersphere (Figure 52B). Temporal stability or consistency of the evolving population was then quantified by the dot product of two time-shifted unity length vectors:

$$S(t) = \widehat{\mathbf{R}}(t-\tau) \cdot \widehat{\mathbf{R}}(t+\tau)$$

The stability metric ($S(t)$) tracks the running similarity of the normalized trajectory separated in time by 2τ msec. Crucially, the normalization constrains $S(t)$ between 0 and 1. Thus, if $S(t) \to 1$, the population activity is considered stable since the relative contribution of each neuron is consistent. If $S(t) \ll 1$, the population activity is deemed unstable because the relative contribution is variable. If the neural trajectory is not normalized, the dot product quantifies similarity across the vectors' magnitude and direction. It roughly mimics the quadratic of the firing rate (exactly so for τ=0). When the vector direction remains constant, the dot product yields no additional information than that already present in the firing rate. In contrast, the normalization scales the neural trajectory so it always has unity magnitude. It neither alters the relative contributions of the neurons nor compromises the vector direction. The dot product therefore performs an unbiased evaluation of stability based only on vector direction, and is an estimate of the fidelity of the population code modulo a multiplicative gain factor. Intuitively, the stability measure is analogous to the correlation between the neurons' activities at two different time points. We chose the dot product, however, because of its interpretability as a measure of pattern similarity in n-dimensional activity space.

Since temporal stability is defined over the whole population and gives rise to a singular measure at a given time, it is not directly amenable to statistical significance analysis. Therefore, we took an alternative approach in order to ensure that the observed stability profiles were contingent on the recorded activity. We shuffled the population activity in two different ways and bootstrapped the stability across multiple shuffles for each type of shuffle (Figure A3). First, we shuffled the activity by randomly assigning each vector in the set of population vectors across time to different time points. This temporal shuffle retains each neuron's identity in the population but scrambles the structure across time. Next, we randomly shuffled the neuron's activities at each time point. This shuffle retains the average firing rate at each instant but removes any temporal

correlation in the firing rate across neurons. We performed multiple shuffles of each kind and computed temporal stability for each instance to obtain bounds for the stability profile in the original data.

5.4.4 Leaky accumulator with facilitation (LAF) model

We developed a computational model to demonstrate how the temporal structure of population activity could be used by neural circuits to gate decisions such as movement initiation. The core idea behind the model is the hypothesis that signal integration is stronger when the temporal structure in input activity is stable, which allows the decoder neuron to reach threshold during the high frequency burst that triggers the movement. We constructed an accumulator as an abstraction of a neuron receiving population inputs through its network of dendrites. The total synaptic current ($I(t)$) and the firing rate ($v(t)$) of the accumulator were defined as

$$\tau_s \dot{I} = -I + \sum_1^n w_i u_i$$

and

$$v = F(I)$$

where u_i and w_i are the instantaneous firing rate and synaptic gain of the i^{th} input neuron, and τ_s is the time constant of the synaptic current. The output firing rate of the single decoder neuron ($v(t)$) can be described by a standard monotonic function (e.g., linear or sigmoid) applied to the net current. The family of stochastic, leaky accumulator models that integrate the firing rate of neurons toward a threshold criterion has been commonly used to explain reaction time, decision making, and perception (*refs*). We used the following heuristic in order to extend this framework to incorporate temporal structure. In order to assess stability, the decoder neuron must keep track of the short term history of the population activity, use this "memory" to evaluate stability, and respond selectively when the activity pattern is deemed stable over some time scale. We implemented these requirements by introducing short-term plasticity in the form of facilitatory connections from the input population onto the output unit (accumulator). The change in connection strength or gain of each neuron on the decoder neuron (w_i) can be defined by a differential equation that incorporates a Hebbian-like learning rule and leak current:

$$\tau_w \dot{w}_i = -w_i + \frac{u_i v}{wnf_i}$$

The Hebbian-learning component, ($u_i v$) describes the change in weight coupled to the firing rates of the i^{th} pre-synaptic neuron and the post-synaptic accumulator neuron. τ_w is the time constant of the weight parameter. Since this version of the model contains only excitatory connections, the weight parameter can exhibit unbounded accumulation, which can be controlled by incorporating normalization. We normalized the Hebbian-learning component ($u_i v$) with the contribution of that neuron to the total resource pool. The resource pool available for facilitation at any given time was defined as the sum of the Hebbian-learning component across all input units ($\sum_1^n u_i v$). The contribution of a neuron to the output rate v then determines its contribution to the overall pool, giving rise to the weight normalization factor:

$$wnf_i = \sum_1^n u_i w_i u_i$$

For the model simulation, we used visuomovement neuron population activity in SC as inputs. We created two sets of input snippets from the visual and premotor bursts, each 180 ms in length. For the visual burst, we used the activity upto the peak in the population response. For the premotor burst, we used activity upto 35 ms before saccade onset, when the response magnitude reached the same level as the peak of the visual burst. In order to control for the effect of average firing rate on the accumulation, we mean-matched the input profiles as follows. We divided the activations in the premotor input set at each instant by the mean activation of the visual inputs at the corresponding instant. That is,

$$u_i^{pre}(t)_{mm} = \frac{u_i^{pre}(t)}{\frac{1}{n}\sum_1^n u_i^{vis}(t)}$$

where $u_i^{vis}(t)$ and $u_i^{pre}(t)$ are the activity of neuron i at time t in the visual and premotor input sets, respectively, and $u_i^{pre}(t)_{mm}$ is the premotor input instantaneously rescaled to match the mean of the visual input. This ensured that we isolated the effect of temporal structure, independent of firing rate, on the model's response. We also quantified the ability of the synaptic weights to track the inputs by computing the Pearson's correlation between the weight and input vectors at each time point. To quantify the accumulator's ability to discriminate temporal pattern in population input, we computed a pattern sensitivity index (psi) as

$$\mathrm{psi}(t) = \frac{v^{pre}(t) - v^{vis}(t)}{v^{pre}(t) + v^{vis}(t)}$$

where v(t) is the output of the accumulation for the corresponding inputs.

5.5 RESULTS

Zooming in on the activity profiles of individual neurons at the time of the visual and premotor bursts (Figure 52A, insets) sheds light on a possible solution to the puzzle of why the visual response in visuomovement neurons does not trigger a saccade. Qualitatively, the neurons' firing profiles seem to be unstructured in the visual burst, while they rise more coherently in the premotor burst. We hypothesized that the degree to which the relative levels of activation across the population is preserved over the course of a burst controls saccade initiation. To quantify this, we created a pseudo-population of our single-unit recordings and constructed a trajectory of population activity over time. We normalized the population vector at each time point, which constrains it to a unit hypersphere in state space. This step factors out global changes in firing across the population and focuses on the relative activity levels. We then computed the dot product between two of these unit vectors separated in time (20 ms here) - we call this the temporal stability of the population code. This procedure is schematized in Figure 52B (for details, see Methods and Figure A1). Figure 52C shows the evolution of temporal stability of the SC and FEF populations (solid and dotted traces, respectively, inset shows individual animals). In both regions, the population code is relatively unstable during the visual burst and remains stable during the premotor burst. This was true across a range of separation times between the population vectors (Figure A2). Thus, the stability across the population of visuomovement neurons seems to function as a switch that prevents movement initiation at an undesirable time (visual epoch), allowing the animal to successfully perform the task at hand (while presumably also allowing the activity to be routed to other brain areas for further processing). Once cued to execute a gaze shift, the activity in the same population rises in a stable manner allowing saccade initiation. Scrambling the population code by either shuffling the activations of individual neurons at each time point or the temporal sequence of population activity drastically alters the stability profile (Figure A3), further suggesting that the above result is due to the specific temporal structure in the recorded population.

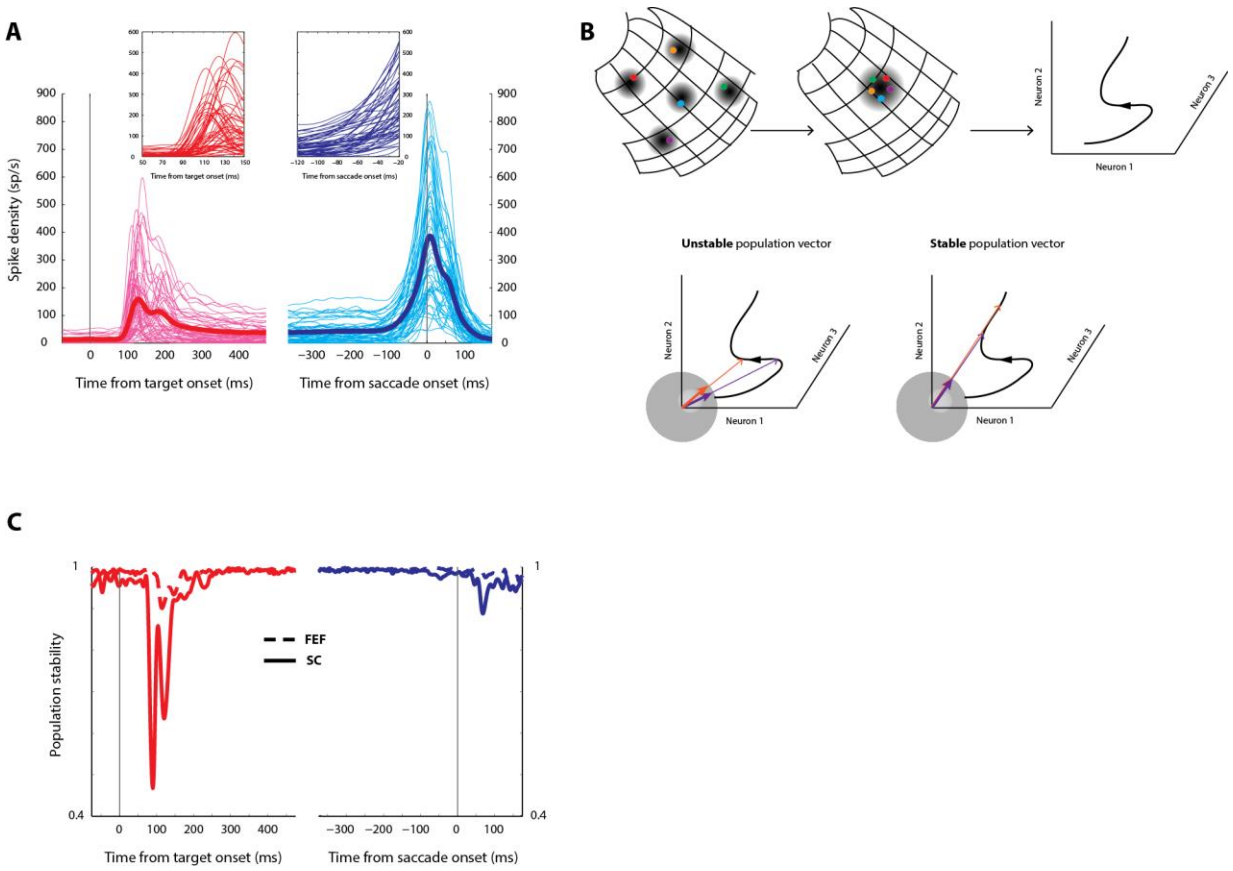

Figure 52. Population activity is temporally unstable in the visual burst and stable in the motor burst

A. Activity profiles of individual SC neurons (thin lines) with the population mean (thick line) overlaid. The zoomed in insets highlight the temporal structure of the two bursts. **B**. Inferring population dynamics from single-unit recordings. The top row depicts the procedure of combining recordings made in different active populations on the SC map (left) into one pseudo-population (middle) that is active for any saccade, enabling the construction of an expected population trajectory (right) in neural state space. The bottom row shows examples of stable and unstable parts of the trajectory. Note that in both cases, the two population vectors (thin arrows in orange and purple) are separated by a similar length of time. In the unstable case, the normalized vectors (thick arrows) move around on the surface of the hypersphere, whereas in the stable case, they stay pointed roughly in the same direction. **C**. Temporal stability is computed as a dot product between two population vectors separated by tau=20 ms for SC and FEF. Line types as in Figure 51. Population stability shows a drastic reduction at the time of the visual burst, but is stable before and during the onset of the movement.

Next, we explored the robustness of the temporal stability hypothesis. Neurons in the rostral SC are tonically active during fixation (Munoz and Wurtz, 1993a) and burst for small saccades and microsaccades (Hafed et al., 2009), but they reduce their activity during larger movements (Munoz and Wurtz, 1993a) (Figure 53A,B). Importantly, they also project to the saccade burst generator (Moschovakis et al., 1996). Assuming population stability sets part of the input drive to downstream structures, we hypothesized that the temporal structure in rostral SC neurons must decrease during large saccades but remain elevated during fixation, even when a

visual burst occurs in other parts of SC and FEF. In other words, we expected the evolution of temporal stability across rostral SC neurons to be the inverse of what occurs in caudal SC - Figure 53C confirms the prediction. Since rostral SC is also known to play a causal role in the generation of microsaccades, we considered whether the population of rostral SC neurons that burst for microsaccades had a stable temporal structure. This was indeed the case, as seen in Figure A4. Thus, the population activity of neurons in rostral SC, including its temporal structure, supplements the pattern in other parts of SC in suppressing and initiating movements. We also examined the stability of the population in the ipsilateral SC with a response field opposite the location of the saccade target (Figure A5). The results were similar to the profile observed for rostral SC, consistent with the idea of reducing the drive from competing neuronal populations. Moreover, pooling neurons in SC and FEF to form a combined population (since they all project to saccade-generating structures) did not impact the main result (Figure A6).

Note that the delayed saccade task is designed to temporally separate the visual and premotor bursts. In the gap task (Figure 53D), in contrast, the two bursts often overlap and saccade reaction times are reduced ("gap effect" (Dorris and Munoz, 1995)); for the subset of movements with the shortest reaction times, the visual and motor bursts merge (Dorris and Munoz, 1995; Edelman and Keller, 1996) (Figure 53E). As another critical test, we hypothesized that if temporal stability is crucial to movement initiation, the visual burst during the shortest latency saccade trials should have a stable temporal structure. This was indeed the case (Figure 53F), providing additional evidence in support of the stability hypothesis.

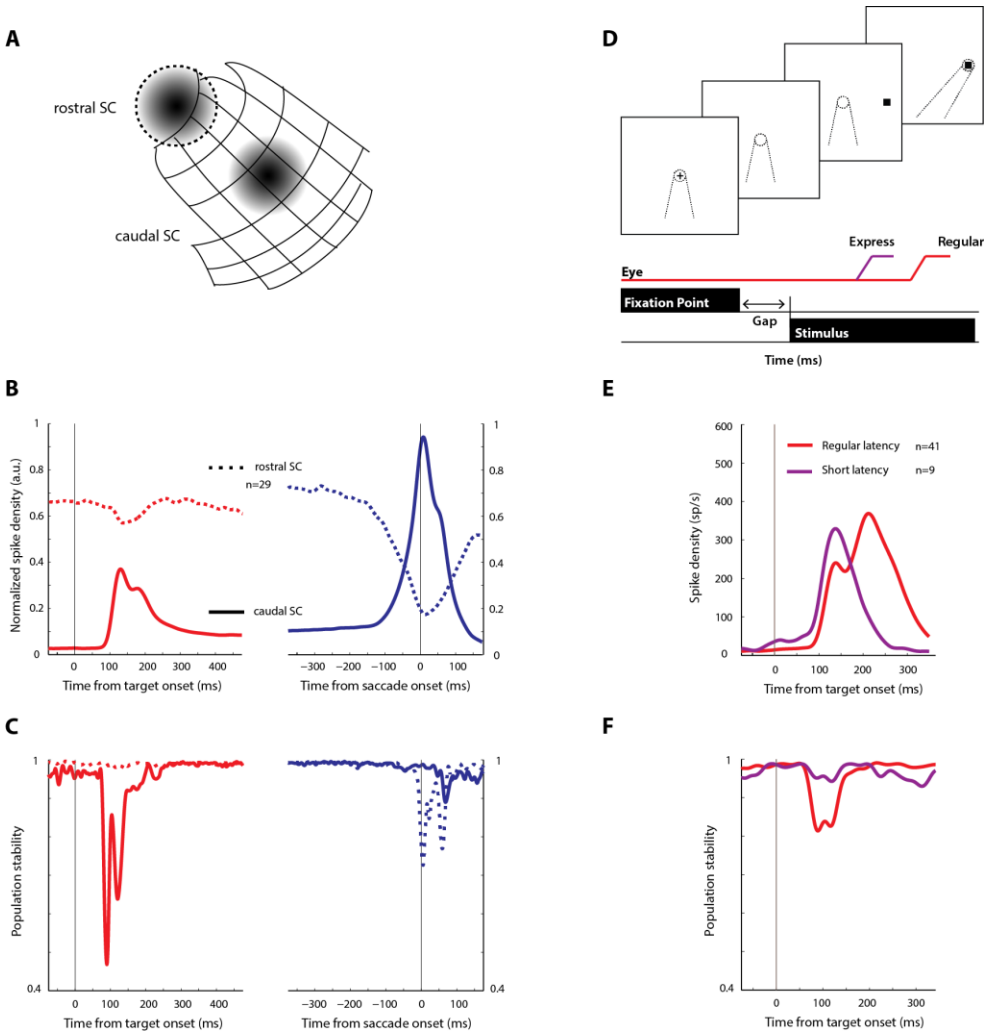

Figure 53. The stability hypothesis is validated for other neural populations and tasks

A. Schematic of the SC depicting active populations in the rostral (dotted boundary) and caudal SC (plain). **B** and **C**. Normalized activity and temporal stability of the rostral SC population (dotted) with the caudal population shown for comparison. **D**. Timeline of the gap task with a schematic of regular (red) and express (purple) saccades. **E**. Mean population activity in caudal SC during the gap task for regular and express-like short latency saccades (colours as in D). Since this is a reactive task, saccade onset is separated from target onset only by 100-200 ms. Data are therefore only shown aligned on target onset for the sake of simplicity. Note the single peak in the burst for the short latency saccades. **F**. Temporal stability during regular and short latency saccades.

Having established that the temporal structures of visual and premotor bursts differ, we sought to identify a mechanism by which downstream neurons could discriminate between stable and unstable population codes. The mere presence of a differential signal does not imply that the brain utilizes that signal in a meaningful manner. In order to provide a framework within which the signal could be used to control behavior, we constructed a rate-based accumulator model that receives input from visuomovement neurons. The accumulator is an abstraction of a neuron receiving population inputs through its network of dendrites (Figure 54A). For simplicity, we

assumed that the output firing rate of the neuron is a monotonically increasing function of the net post-synaptic current (Rauch et al., 2003). How might the population pattern be converted into a signal that initiates a movement only when the pattern being decoded is stable? Heuristically speaking, the decoder should keep track of the short-term history of population activity, use this "memory" to evaluate stability, and respond selectively when the activity pattern is deemed stable over some time scale. We incorporated these requirements in our model by introducing short-term synaptic plasticity in the form of facilitatory connections (Borst and Sakmann, 1998; Tsujimoto et al., 2002) from the input population to the output unit (accumulator), with a Hebbian-like learning rule and leak current. Since the network only contained excitatory connections, we also included weight normalization to prevent unbounded accumulation in the output (Online Methods and Supplementary Info).

We then tested whether the leaky accumulator with facilitation (LAF) model could exploit temporal structure by feeding it population activity from visuomovement neurons (only inputs from SC shown here, Figure 54B1). We focused on two snippets of activity - the 180 ms period ending at the peak of the visual response, and the 180 ms period terminating 35 ms before saccade onset (when the mean activity in the premotor burst reached a level similar to the peak visual response). To discount effects of the mean firing rate, we matched the population means at each time point by scaling the population vector of one set of inputs by the running ratio of their respective means (Figure 54B2). Note that the temporal structure is preserved in each signal (Figure 54B3). The accumulator increased its activity at a faster rate in response to the stable pattern compared to the unstable pattern (Figure 54B4), suggesting that the network was able to discriminate between two types of population patterns even though the net input drive was identical.

How critical is facilitation to this function? Like the inputs themselves, the weights also showed greater fluctuation during the visual burst (around t=140 ms in red traces, Figure 54B5). Consistent with our heuristic, the weights tracked the input rates when the input was stable, but this correlation dropped away when the input was unstable (Figure 54B6). Note that the shape of the correlation profile is not unlike the stability profile in Figure 54B3. Therefore, facilitation allows the weights to retain the memory of a pattern over the time scale of tens of milliseconds, but the memory can be scrambled by a fluctuating input pattern. We quantified the network's

ability to discriminate between the two patterns with a pattern sensitivity index (Figure 54B7). The results here suggest that the accumulator is driven more strongly by the stable premotor burst, even when other features of population activity remain the same.

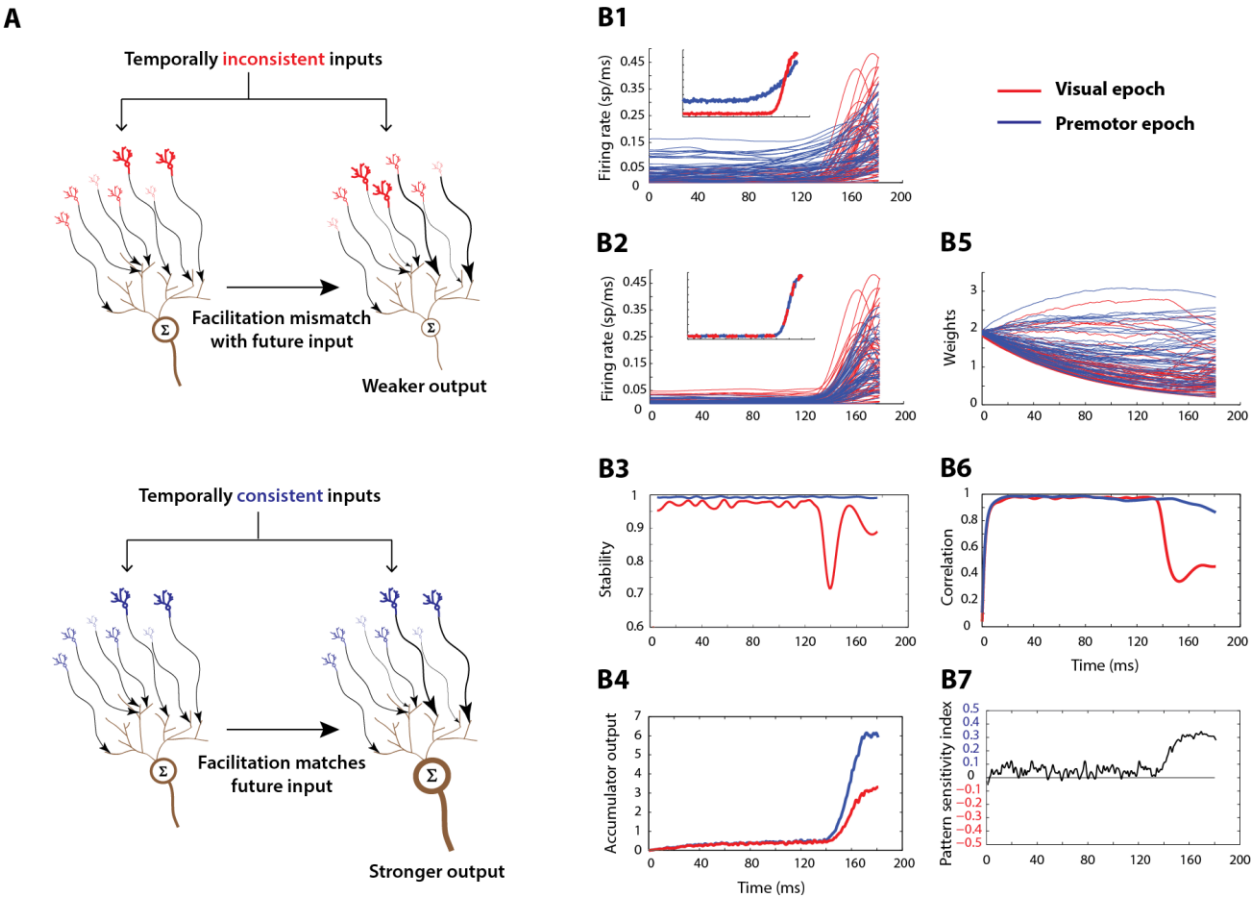

Figure 54. Leaky accumulator with facilitation model can discriminate population temporal structure

A. Schematic of the LAF model. The model incorporates a neuron that receives population inputs through its dendritic network and integrates them to produce an output. Each input neuron's activation level is represented by its size and thickness, and its connection strength to the accumulator is represented by the thickness of the arrow. The synaptic weights can grow over time due to short-term facilitation. Two time points are shown for illustration (left – arbitrary initial point, right – subsequent time point). Top – Unstable inputs lead to a scenario where the facilitated synapses that are no longer driven by the strong inputs that created them, whereas, bottom – stable inputs lead to matched strong input and synapses, resulting in stronger accumulation. **B**. Inputs to the model from SC visuomovement neurons. 1 - Raw unstable (red) and stable (blue) input profiles (inset – population means). The two sets of inputs are 180 ms snippets taken from the visual and premotor bursts, respectively, in the spike density profiles shown in Figure 51. 2 - Mean-matched input profiles and population means. 3 - Temporal stability of the two mean-matched populations. 4 - Output of the LAF accumulator in response to the stable and unstable model inputs. 5 - Evolution of synaptic weights for the two conditions. 6 - Correlation between instantaneous weights and input rates for the two conditions. 7 – Pattern sensitivity index of the model's ability to discriminate between the two types of inputs. Values in the top half of the plot indicate higher sensitivity (faster accumulation) to stable population input.

5.6 DISCUSSION

Neurons in premotor structures are constantly bombarded with information from thousands of presynaptic neurons that are active during sensorimotor processing. It is unclear how activity relevant for movement initiation is discriminated from activity related to other processing. A canonical model of movement initiation, especially in the oculomotor system, is threshold-based gating (Hanes and Schall, 1996). Current knowledge points to a role by the omnipause neurons (OPNs) in defining the threshold and controlling saccade initiation (Hanes and Schall, 1996; Everling et al., 1998). However, thresholds vary across behavioural paradigms (Jantz et al., 2013), raising the question of how the threshold is set in a particular condition. Furthermore, evidence that the threshold changes during the course of a trial purely based on OPN activity is limited (Everling et al., 1998). A related explanation is that a movement is initiated only when movement-related neurons are active, but most neurons in sensorimotor structures likely span a continuum between having visuomovement activity to pure movement activity (Wurtz et al., 2001). Other models posit that a movement is initiated when neural activity traverses certain optimal regions of the population state space (Churchland et al., 2012) and is inhibited otherwise (Kaufman et al., 2014). While the optimal subspace and nullspace hypotheses are certainly appealing, a mechanistic understanding of why certain patterns of activity translate to movement commands is lacking. We reason that for a decoder downstream, it is important to ensure the stability - or consistency over time - of the input code before processing it to influence output. Population activity that creates a high firing rate drive but is inconsistent should be prevented from triggering a movement. This requirement is critical especially in the case of ballistic movements such as saccades, where the ability to reverse the decision once the movement has been initiated is limited.

Our findings are closely related to the premotor theory of attention (Rizzolatti et al., 1987) and offer a way to reconcile the attention-or-intention debate. They could also account for the mirror-like activity recorded during both action observation and execution in neurons known to project directly to motoneurons in the spinal cord (Vigneswaran et al., 2013). In both cases, it is unclear how the same neuronal population represents two distinct signals that serve different functional roles. Our results suggest that this multiplexing ability may be provided by the distinct temporal structures of population activity patterns. Previous studies have looked at the role of precise synchronization in the timing of incoming spikes in the transmission of information and

the efficacy of driving the recipient neuron (Carter et al., 2007; Branco et al., 2010). However, input firing rates may vary greatly across the population, limiting the ability to compare to spike times. Our study proposes a mean-field equivalent to the spike-based temporal or correlation code (Softky, 1994; Harvey et al., 2013) by looking at the temporal structure of population firing rates, thus tying together the notion of rate, temporal, and population codes. We suggest that temporal structure of population activity is a critical to understanding movement generation as well as, more broadly, all neuronal communication and its relationship to behaviour.

6.0 DISCUSSION AND CONCLUSIONS

6.1 THE SENSORIMOTOR MULTIPLEXING PROBLEM

The primate brain has evolved several strategies to successfully navigate the complex environment we are embedded in. An efficient way to parse important aspects of the visual environment is by shifting the line of sight (or gaze) to different locations in the scene. Gaze shifts involve a sensory-to-motor transformation of the visual input coming in from the retina to a movement command instructing the eye (and, generally speaking, the head) to rotate by the appropriate amount to land the image of the new stimulus on the fovea. This transformation is implemented by a distributed network of brain regions that together comprise the gaze control network (Wurtz et al., 2001). In fact, there exist neurons in this network that are directly equipped to implement this transformation. These so-called visuomovement neurons produce a burst of activity both in response to the onset of a visual stimulus and during an eye movement (or saccade) to that stimulus (Sparks and Mays, 1990; Dorris et al., 1997). Many of these neurons also maintain an elevated level of discharge during the period between the onset of a stimulus and the actual saccade, a period during which the sensorimotor transformation evolves before maturing into the motor command. This low frequency activity has been attributed to a number of cognitive processes, including target selection (Schall and Hanes, 1993; Horwitz and Newsome, 1999; McPeek and Keller, 2002; Carello and Krauzlis, 2004), spatial attention (Goldberg and Wurtz, 1972; Ignashchenkova et al., 2004; Thompson et al., 2005), working memory (Balan and Ferrera, 2003; Zhang and Barash, 2004), decision-making (Newsome et al., 1989; Gold and Shadlen, 2000), and motor preparation (Hanes and Schall, 1996; Kustov and Robinson, 1996; Dorris et al., 1997). Intriguingly, visuomovement neurons in the superior colliculus (SC) also project directly to burst generators in the brainstem that issue the final command to produce a saccade (Segraves, 1992; Rodgers et al., 2006). This poses a conundrum – why does the sensory burst not produce a

movement? The work in this thesis was primarily motivated by the lack of a clear understanding about how downstream structures distinguish between activity related to sensory and cognitive processing and activity related to movement preparation. We call this the sensorimotor multiplexing problem.

6.2 POSSIBLE SOLUTIONS

One possible way to resolve the sensorimotor multiplexing problem is to invoke the idea of division of labour (Steinmetz and Moore, 2012). Apart from visuomovement neurons, gaze control structures such as SC typically also contain movement-only neurons that also project to the brainstem burst generators (Cohen et al., 2009). These neurons do not, under normal conditions, respond to the onset of a sensory stimulus, only becoming active prior to a saccade. Therefore, it has been suggested that activation of the burst generators to produce a movement requires the combined activity of both visuomovement and movement neurons, and that visuomovement activity is integrated to threshold to generate activity in movement neurons (Schall et al., 2011). However, this explanation just passes upstream the onus of distinguishing movement-related activity from sensory and cognitive processing-related activity. Why, one may ask, do the movement neurons become active when they do – just before a saccade is produced – and not earlier? Moreover, as we show in Chapter 2 of this thesis, movement neurons can also exhibit stimulus-locked activity in some circumstances (Figure 23). This suggests that the visuomovement/movement divide is not rigidly grounded in physiology and that neurons in sensorimotor structures can flexibly modulate their activity depending on the state of the rest of the network.

Another way around the problem is to postulate a threshold-based mechanism for movement initiation. Neurons in the frontal eye fields (FEF) have been shown to reach a fixed level of activation before the onset of a saccade (Hanes and Schall, 1996), suggesting that neural activity needs to reach a threshold before a movement can be initiated. Subsequent studies have shown that the threshold is not fixed for a given neuron and can change depending on task requirements (Jantz et al., 2013). Nevertheless, it is possible that the threshold is implemented at

the level of a neuronal ensemble (Zandbelt et al., 2014). In Chapter 4 of this thesis, we show that saccades can be produced without reaching the purported threshold, at the single neuron or population level, under identical task conditions (Figure 43). Moreover, there exist single neurons in which the visual burst crosses the computed threshold, yet no movement is produced (Jantz et al., 2013, also see Figure 52).

A related idea is downstream gating of sensorimotor activity – the gate is opened only to allow activity from the premotor burst to go through to the brainstem. The omnipause neurons (OPNs) in the paramedian pontine reticular formation are the likely physiological correlate of this proposed gate. These neurons are tonically active when the eyes are stationary and turn off during a saccade (Cohen and Henn, 1972). However, like the division of labour hypothesis presented above, we think this explanation merely shifts the burden of distinguishing sensory and movement-related activity to downstream structures. To wit, how do the OPNs know when to turn off?

Note that the primary reason sensorimotor multiplexing presents a problem is because of the tacit assumption that neural correlates of sensory and movement processing are represented by the instantaneous firing rate of the neurons. While firing rates at the individual neuron or ensemble average level are sufficient to explain a variety of cognitive and behavioural phenomena, the scalar nature of those quantities restricts their extension to higher-dimensional correlates, e.g., multiplexing. Several recent studies have acknowledged this limitation and have explored the idea of coding by latent neural ensembles in multi-dimensional space (Churchland et al., 2012; Mante et al., 2013; Rigotti et al., 2013), partly owing to the increasing use of high-density neural recordings. Studies of movement initiation in the skeletomotor system have used this approach and found that neural activity reaches a specific region of neural state space before undergoing dynamics that produce a limb movement – the "optimal subspace" hypothesis (Afshar et al., 2011; Churchland et al., 2012). Related work has shown that activity pertaining to movement preparation evolves in a region of population space that is orthogonal to the optimal subspace, (i.e., the null space), and this dissociation allows the animal to prepare the movement by incorporating perceptual and cognitive information without risking a premature movement (Kaufman et al., 2014). It remains to be seen whether these concepts are applicable to sensorimotor transformations in the gaze control system.

In Chapter 5, we proposed a novel hypothesis to resolve the sensorimotor multiplexing problem that taps into population dynamics. We showed that the relationship between firing rates across the population of SC visuomovement neurons evolves differently during the stimulus-locked burst compared to the movement-related burst. Population activity structure becomes inconsistent over the tens of milliseconds spanning the visual burst but was maintained during the saccade (Figure 52). We further showed that a simple decoder capable of tracking population input patterns over time is sensitive to the temporal structure of those inputs and can differentially integrate consistent and inconsistent patterns (Figure 54). This idea has intuitive appeal; downstream neurons typically receive inputs from hundreds and thousands of sources, and the ability to resolve temporal dynamics in those inputs affords greater range and flexibility in decoding space. Moreover, it meshes well with studies that have shown that the timing, or temporal sequence, of incoming action potentials can significantly impact dendritic integration and the likelihood of the post-synaptic neuron to generate an action potential (Carter et al., 2007; Branco et al., 2010). Thus, the neural circuitry involved in sensorimotor integration may have evolved to delineate sensory, motor and cognitive signals based on the dynamical structure of population activity induced by those factors.

Figure 61 summarizes the models discussed above by re-organizing them in the framework of multi-dimensional population activity.

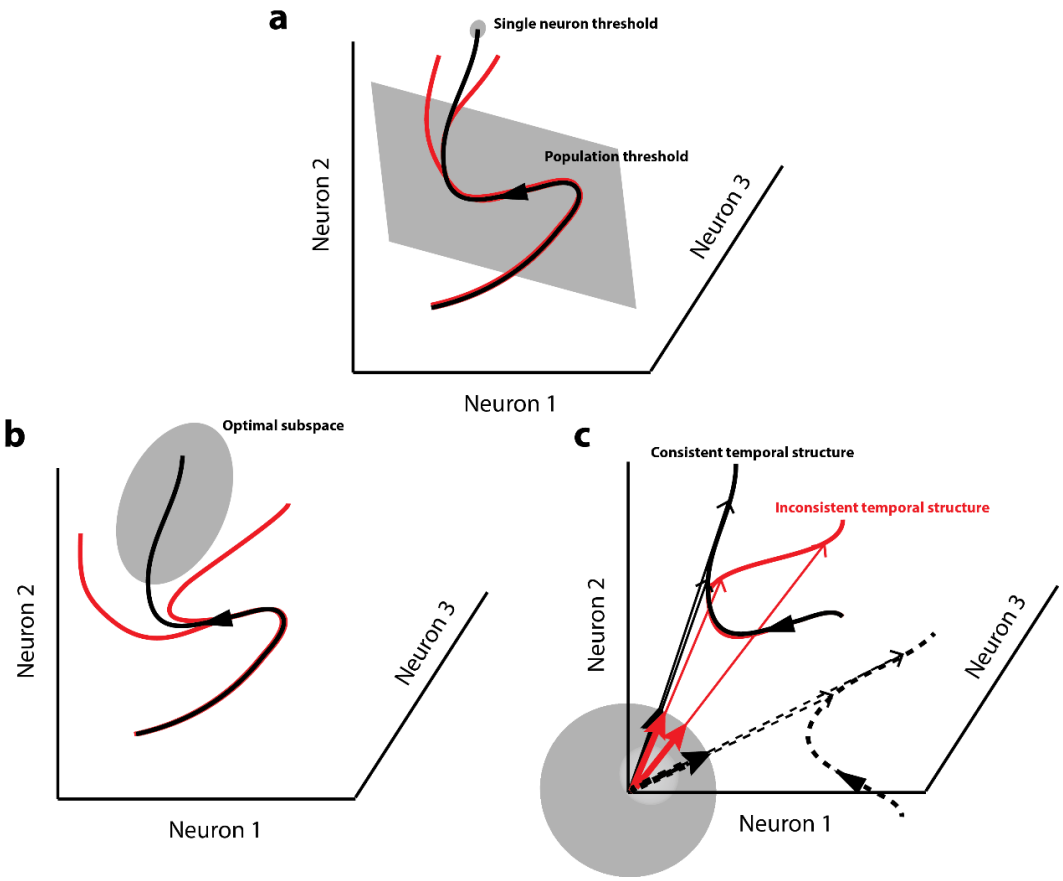

Figure 61. Summary of models of movement preparation

a. Multi-dimensional representation of the threshold hypotheses. Two scenarios are depicted in this figure. One possibility is that the activity of each neuron must rise to a fixed threshold at the same time in order to initiate a movement. Fixed thresholds for each neuron are equivalent to a point (or a small region, represented by the gray sphere) in neural state space. Activity profiles that meet this criterion (e.g., black trace) lead to movement generation while those that don't meet this criterion (red traces) do not result in a movement. Alternatively, activity may need to cross a fixed threshold at the population level (sum of neurons' activities = constant, i.e., an n-1 dimensional hyperplane, gray plane). **b.** The optimal subspace hypothesis offers more leeway by allowing neuronal activity to reach a relatively larger region of population state space (black trace and gray ellipsoid) over a period of time in order to signal the initiation of a movement. Activity trajectories that evolve outside this "optimal subspace" (red traces) do not lead to a movement. **c.** The temporal stability hypothesis suggests that neural activity that is consistent over time in state space, i.e., an activity trajectory that points in the same direction (solid black trace and vectors) is likely to be interpreted as a movement command by a decoder, while activity that is inconsistent (fluctuating directions, red trace and arrows) is not. It does not matter where in state space the activity happens to evolve – a different subpopulation of neurons could be active, but as long as they are pointing in the same direction (dotted black trajectory and vectors), it will lead to a movement. The gray sphere around the origin represents a unit hypersphere for visualization of back-projected unit vectors.

6.3 SENSORIMOTOR MULTIPLEXING AND THE PREMOTOR THEORY OF ATTENTION

More broadly, the sensorimotor multiplexing problem can be reframed in the context of the premotor theory of attention, which posits that spatial attention is instantiated in the same networks that generate movements (Rizzolatti et al., 1987). In other words, low frequency (or "sub-threshold") activity in neurons that generate saccades manifests as an observable phenomenon like attention. In this interpretation, spatial attention and the intention to shift gaze to the attentional locus are equivalent processes (Hoffman and Subramaniam, 1995; Kustov and Robinson, 1996). Studies with cleverly designed behavioural paradigms have made observations that are consistent with the premotor theory (Kustov and Robinson, 1996; Gold and Shadlen, 2000), as well as those supporting the notion of serial transformations (Thompson et al., 1996; Horwitz and Newsome, 1999). Efforts to delineate attentional and movement processing by means of causal manipulations have also produced mixed results, with some studies supporting disjoint processing (Juan et al., 2004) and others supporting parallel processing (Moore and Armstrong, 2003; Katnani and Gandhi, 2013).

The premotor theory is ethologically elegant – it reduces the need to invoke modularity in neural hardware to explain the presence of high level cognitive functions and simplifies the question of how such functions could have emerged to begin with. Consider that evolutionarily primitive animals have a limited repertoire of stimulus-response mappings and short reaction times, resulting in impulsive and predictable behaviour. It may have been advantageous to delay pre-potent actions and increase response flexibility. One way to achieve this efficiently would be to use existing motor circuitry hardware and alter the coding scheme so that a volley of activation caused by sensory inputs no longer directly triggers a movement but is maintained as low-frequency activity in the premotor network and routed to other brain areas for further processing (creating perception and cognition). It is possible that this dissociation is achieved by discrimination boundaries in population activity space (optimal vs null subspace), or in dynamical features of evolving activity (consistency of temporal structure).

An advantage of such a mechanism is the ability to produce movements when needed, since the information about the action plan is immediately and concurrently available to the motor

circuitry (albeit along different dimensions than what is needed to initiate a movement), while allowing the activity to be used for other purposes. For example, apart from the functions mentioned earlier, it has been suggested that low-frequency activity in premotor networks supports complex cognitive functions such as mental imagery (Kosslyn et al., 2001) and movement rehearsal (Cisek and Kalaska, 2004). Mental imagery of an action is externally induced during action observation (watching another perform the action), a function that is mediated by neurons in premotor cortex that are also active during action execution (Gallese et al., 1996). Just like sensorimotor neurons in the gaze control network, these so-called mirror neurons project directly to the spinal cord (Kraskov et al., 2014), and have been proposed to play a role in action suppression (Kraskov et al., 2009; Vigneswaran et al., 2013). However, the projections are mostly excitatory, suggesting that the mechanism of action suppression may play out at the level of population dynamics, instead of through a simple weighted summation of projection signs.

6.4 CONCLUDING REMARKS

The findings in this thesis have the potential to significantly impact our understanding of sensorimotor control in the healthy and dysfunctional brain, inform future directions in the scientific study of sensorimotor systems and improve the ability to convert that knowledge to neural engineering applications, such as brain computer interfaces and therapeutic microstimulation devices. The finding of a motor potential in sensory and low-frequency preparatory activity may help us understand how sensorimotor integration is affected in neuropsychiatric disorders like ADHD (Munoz et al., 2003) and schizophrenia (Braff and Geyer, 1990), where the ability to suppress impulsive movements is compromised (Reuter and Kathmann, 2004). Causal experiments using spatially and temporally patterned micro- or optogenetic stimulation (Ferenczi and Deisseroth, 2016) to direct network dynamics into arbitrary regions of activity space may shed light on the precise mechanisms of movement initiation, and differentiate between sensory, cognitive and movement-related processing at the population level, helping to resolve the attention-intention debate in the near future. Incorporating temporal patterning in decoding algorithms for neuroprosthetic and simulated reality devices may improve their performance by providing the ability to better discriminate activity related to movement planning and higher cognitive processes that may be instantiated in the same networks.

APPENDIX A

SUPPLEMENTARY MATERIAL FOR CHAPTER 5

This appendix contains additional figures and derivations referenced in Chapter 5 of this document.

A.1 SUPPLEMENTARY FIGURES

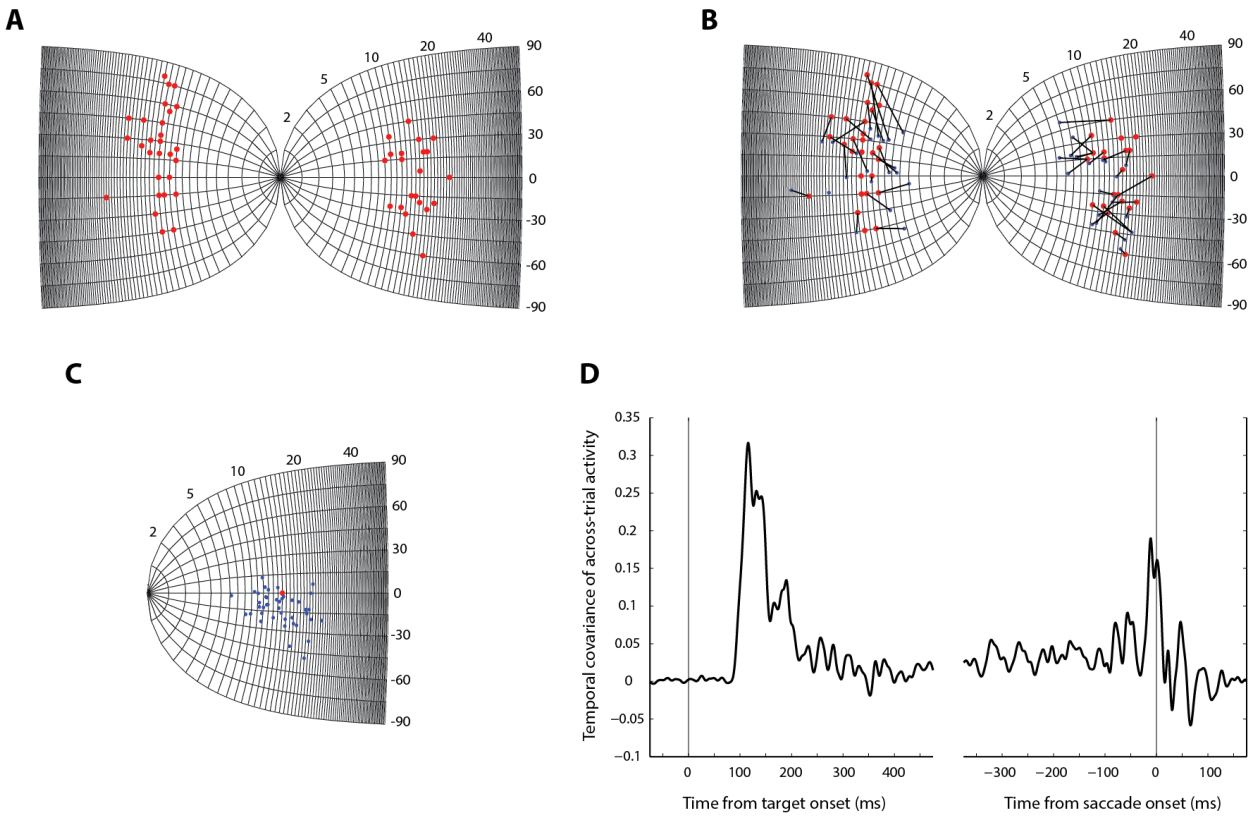

Figure A1. Inferring population dynamics from single-unit recordings

A. Point images (red) of target locations on the SC map across all experimental sessions. The locations on the SC were computed using known transformations between visual space and SC tissue coordinates (Online Methods). **B**. Same as in A, with the endpoint of the stimulation-evoked saccade vector at the recording site shown in blue. The stimulation vector provides a proxy for the RF center of the recorded neuron. Neuron-target pairs (blue-red) from individual sessions are connected using black lines. **C**. The active pseudo-population on the SC map. The red locations from A and B were referenced to an arbitrary location on the SC map (here, R=15, theta = 0) and the blue locations appropriately translated. Points from both colliculi have been shown on a single SC for the sake of clarity. **D**. Covariance of activity over time and across trials averaged across all visuomovement neurons in SC aligned on target (left) and saccade (right) onsets. For a derivation of how using the average activity for each neuron sets limits on the expected stability on single trials, see Supplementary Material.

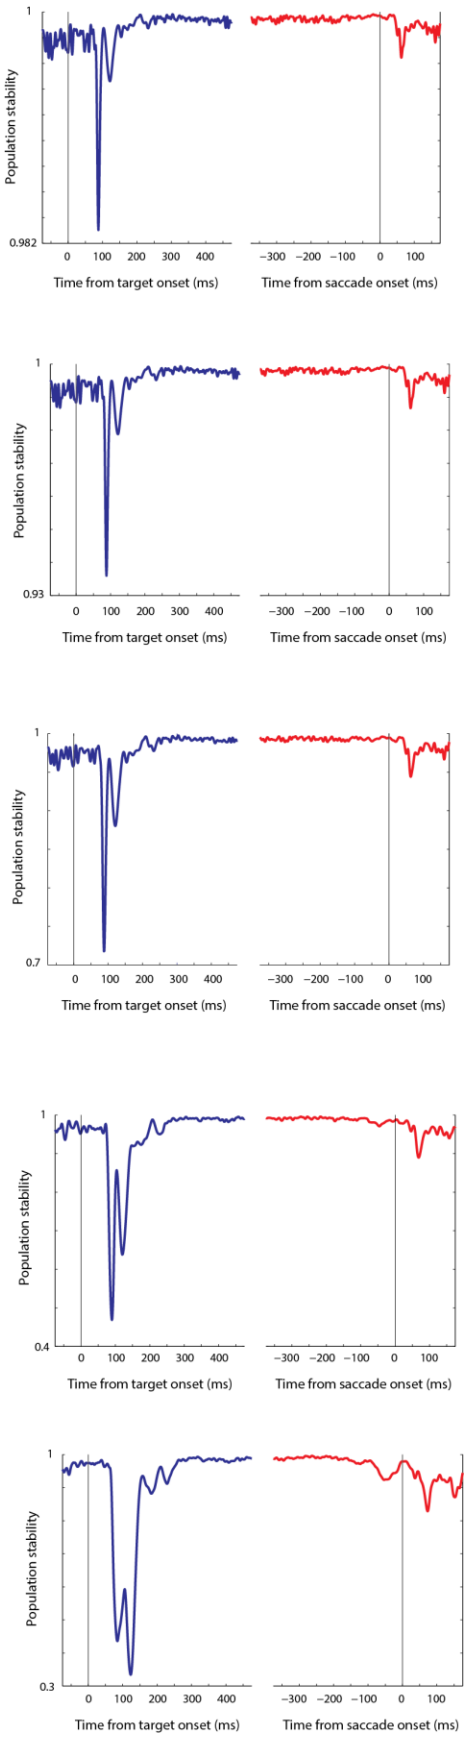

Figure A2. Temporal stability as a function of τ

In order to ensure that the choice of τ did not have an undue effect on our results, we computed stability for several values of τ. The panels show temporal stability profiles for SC as in Figure 2C, for (from top to bottom), τ = 1, 2.5, 5, 10, and 20, respectively. As seen, the absolute magnitude of stability was inversely related to τ - a consequence of the smooth and continuous nature of the trajectory. In other words, since the state of the neural population evolves smoothly, it must traverse intermediate states in order to move from one state to another, resulting in greater similarity between state vectors close together in time than those further apart. However, the relative shape of stability profile was largely preserved across τ (except for some deviation for values of τ as high as 40 ms, possibly a consequence of comparing population vectors from different epochs). Thus, the relative instability during the visual epoch and stability during the premotor epoch persists across all time separations.

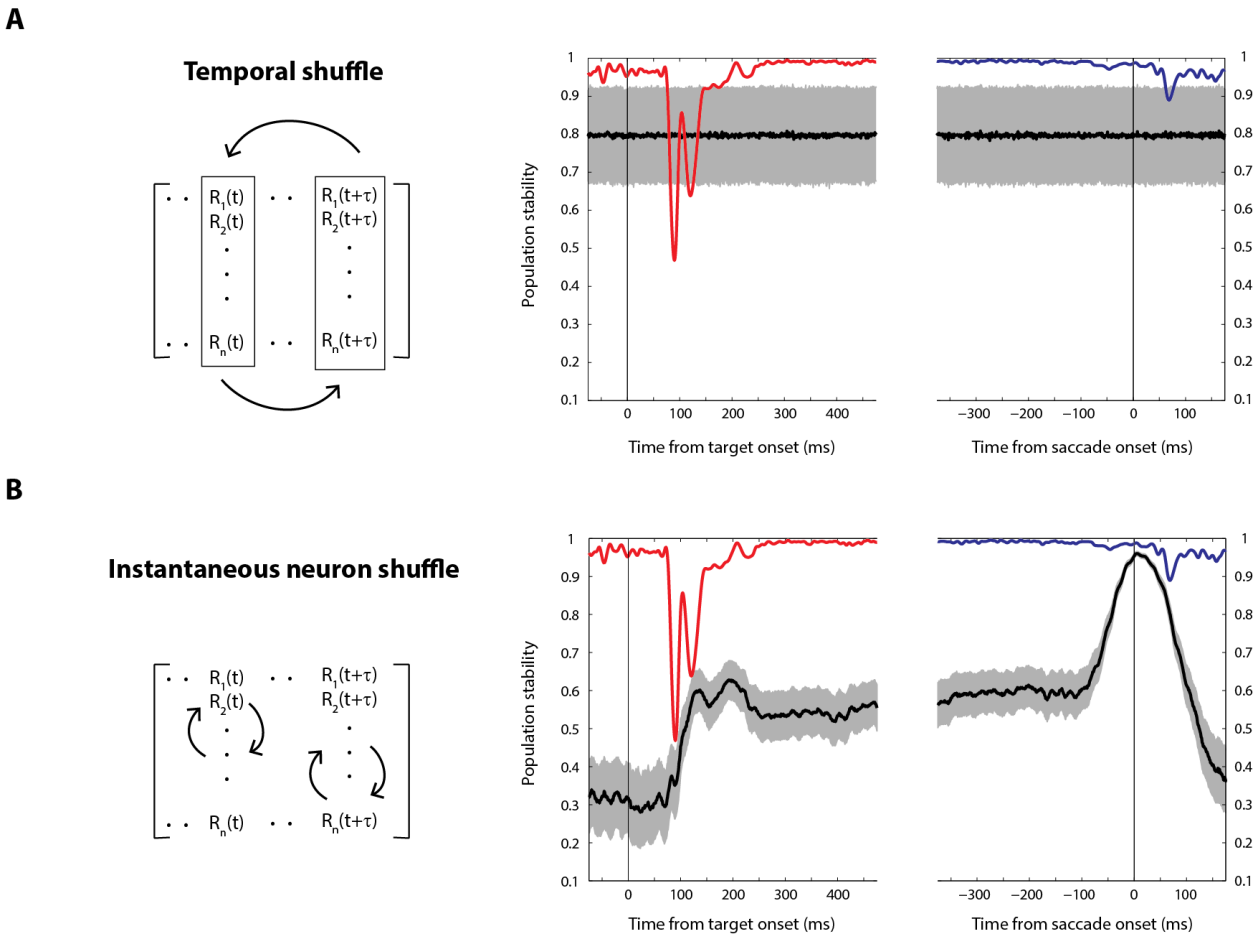

Figure A3. Temporal stability profile is specific to the recorded population structure

A. Left - The population code was altered by randomly shuffling the population vectors between different time points, i.e., the activity vectors are taken from the population trajectory in the original data but in a different temporal sequence. Right - The black trace is the mean temporal stability across 1000 shuffles, and the gray region is the standard deviation across these samples. The stability profile flattens at what is effectively the average value of the dot product between any two vectors in the trajectory. This bootstrap approach indicates that the observed stability profile (red and blue traces) is specific to the true population dynamics. **B**. Left - The population code was scrambled by randomly shuffling the neuron's activities at each time point. This shuffle preserves the average population activity at each instant while removing any ordering between individual neurons. Right - The observed stability profile (red and blue traces) is well outside the distribution of profiles obtained across 1000 shuffles (mean – black trace, standard deviation – gray region).

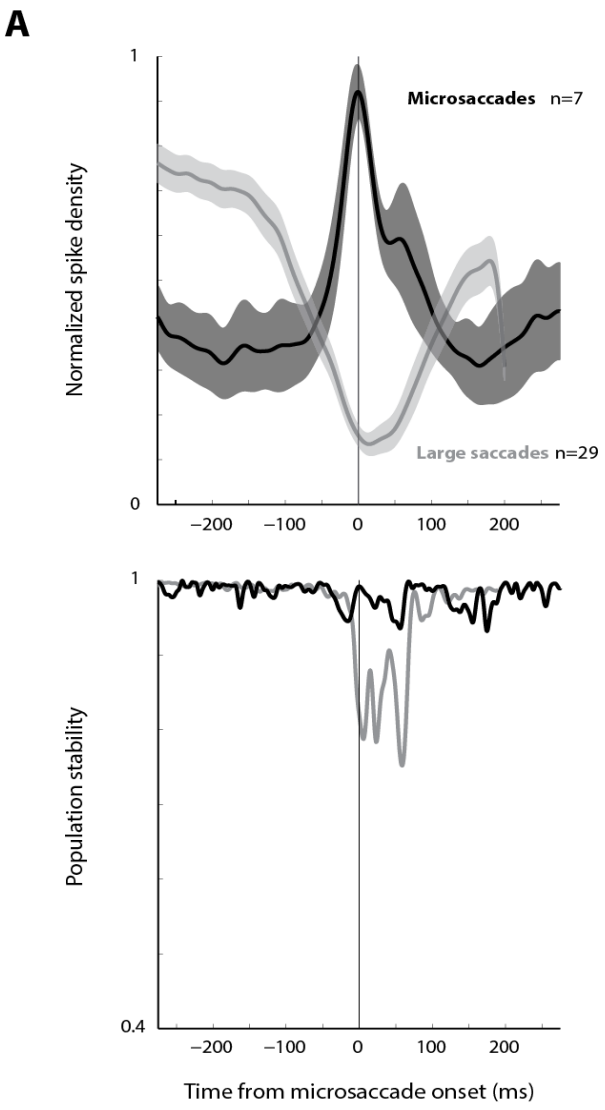

Figure A4. Population activity of rostral SC neurons is stable during microsaccades

Top – Normalized population activity of microsaccade-related neurons in rostral SC (black). Note the strong burst compared to the suppression in fixation-related neurons in the same region for large saccades (gray). Bottom – Temporal stability of the rostral SC population for microsaccades shown in comparison to large saccades. Neurons which burst for microsaccades remain stable during these fixational eye movements, consistent with the hypothesis that both an increase in firing rate and high stability is required for movement generation.

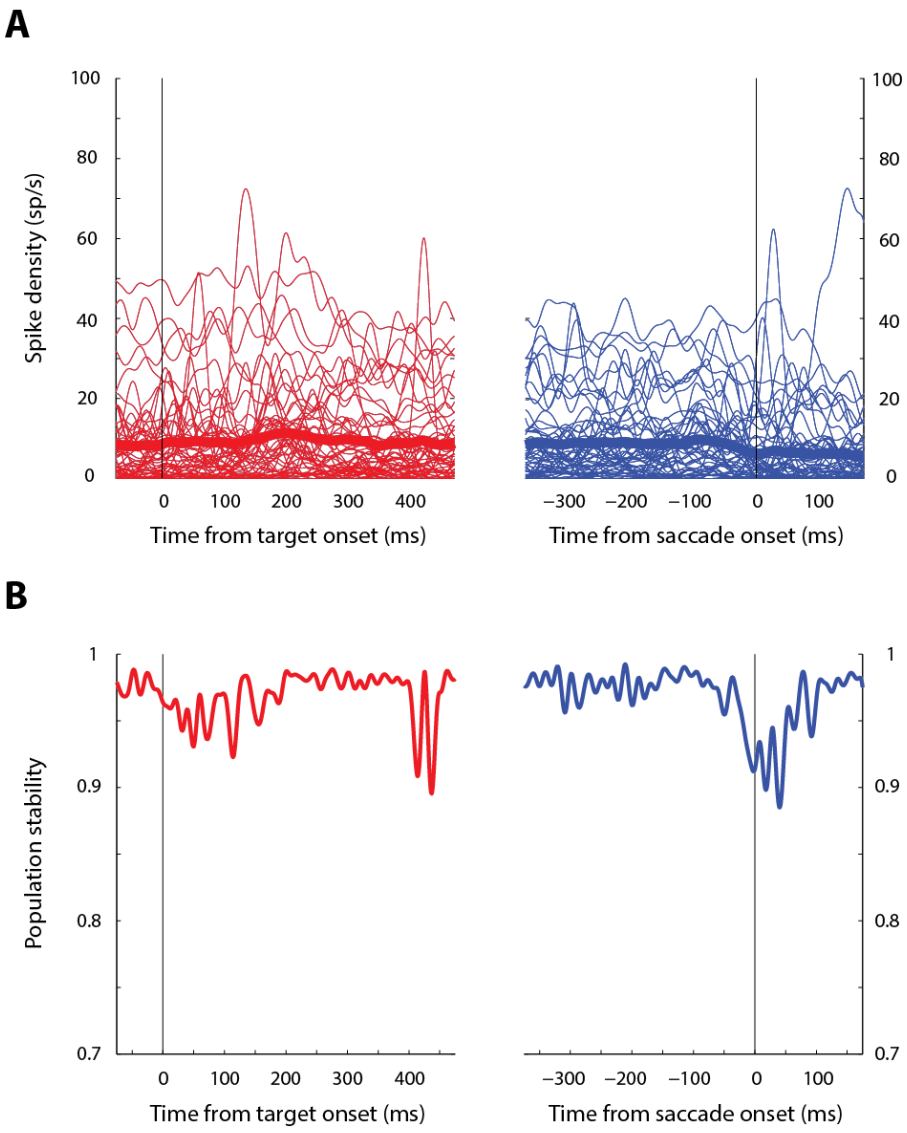

Figure A5. Characteristics of the population in ipsilateral SC

A. Spike densities of individual neurons (thin traces) and the population average (thick trace) when the target (left panel) and saccade (right panel) were in the ipsilateral hemifield. These neurons show little to no modulation in their firing rate when neurons in the contralateral SC actively produce bursts. **B**. Temporal stability of the ipsilateral population aligned on target (left) and saccade (right) onset. Like the rostral SC (Figure 3), the ipsilateral population is relatively more unstable during the saccade, suggesting that stability may be an actively controlled variable designed to reduce the effective drive from competing neuronal populations during movement generation.

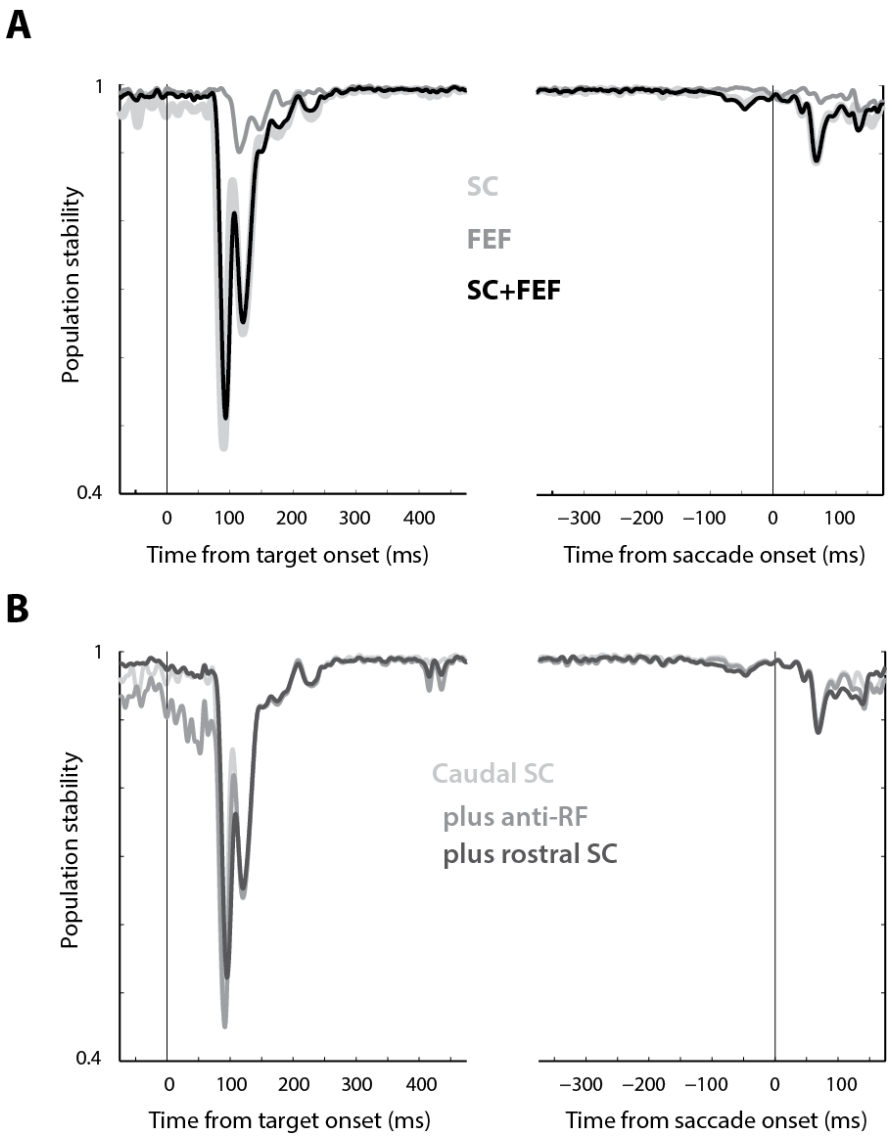

Figure A6. Temporal stability of combined populations

A. Temporal stability aligned on target (left) and saccade (right) onset for neurons constituting the individual "active populations" in caudal SC and FEF (same as Figure 2) in the two shades of gray, and the combined active population in black. Since the burst generators receive input from both sets of neurons, it is important to verify that the characteristics of population temporal structure is preserved for the whole population. **B**. Temporal stability for the active caudal SC population (light gray, same as above and Figure 2), combined caudal SC population (including the ipsilateral SC from Figure A4, medium gray), and all neurons in SC (dark gray). Since the ipsilateral and rostral SC populations contribute little to the firing rate vector, their contribution to the dot product is also low – therefore, any difference between these stability profiles is small.

A.2 INFERRING POPULATION DYNAMICS FROM SINGLE-UNIT RECORDINGS

We previously described in the main text and methods how we reconstructed a pseudo-population from unit recordings. Using knowledge of the RF centre of a given session's local population obtained from microstimulation, and known transformation from visual space to SC tissue coordinates, here we reconstruct the pseudo-population on the SC map.

We first transferred target locations (R_T, θ_T) inside the RF for each neuron onto the SC map using the following transformations (Ottes et al., 1986) –

$x_T = \frac{B_x}{2} \ln \frac{(H+A)^2+V^2}{A^2}$, $y_T = B_y \tan^{-1} \frac{V}{H+A}$, where,

$H = R_T \cos\theta_T$, $= R_T \sin\theta_T$, and $A = 3$, $B_x = 1.4$, $B_y = 1.8$.

The transformed locations (x_T, y_T) on the bilateral SC map are shown in Figure A1, panel A. In order to move these target locations to one pseudo-population, we need to identify where these points reside in the RF of each neuron or local population. Previous studies have demonstrated a strong correspondence between the saccade vector evoked by microstimulation and the centre of the RF of neurons at the site of microstimulation. We used the endpoints of the site-specific microstimulation-evoked saccades as a proxy for the respective RF centres. The transformed endpoints, with their locations relative to the corresponding target locations, are shown in Figure A1, panel B. We picked an arbitrary location in the visual field as the RF centre of the pseudo-population, $(R_c, \theta_c) = (15,0)$. We used the fact that the size of the active SC population is invariant regardless of the encoded saccade to preserve relative distances in tissue space between a site's target location and RF centre, and translated the transformed microstimulation endpoints to the common pseudo-population centre, along with the respective target locations. To construct one active population from sites gleaned from both hemifields (and therefore both colliculi), we reflected the coordinates from one colliculus onto the other. The resulting pseudo-population (Figure A1, panel C) shows a fairly representative sampling of neurons from the pseudo-population, with its extent consistent with a large (25% of SC tissue) active population, albeit one that is biased to one side of the population.

A fundamental assumption behind treating the expected population trajectory as a proxy for an arbitrary single-trial trajectory is that relevant features of single-trial population dynamics are captured by the expected population trajectory. However, it is possible that correlated dynamics between neurons, such as noise correlations on individual trials, may cause measures based on

single-trial trajectories to deviate significantly from the dynamics of the average. We show here that our measure of the structure of the trial-specific population code – temporal stability – is indeed preserved in the expected trajectory.

Let us denote by $\hat{R}(t)$ the population unit vectors comprising the expected trajectory. Let us also denote the population unit vectors making up individual trial trajectories by $\widehat{{}_kR}(t)$, where k is a trial index. The expectation of the inner product stability measure across trials is given by,

$$E\left[\widehat{{}_kR}\left(t-\frac{\tau}{2}\right).\widehat{{}_kR}\left(t+\frac{\tau}{2}\right)\right] =$$

$$E[\sum_i {}_kR_i(t-\frac{\tau}{2})\ {}_kR_i(t+\frac{\tau}{2})] =$$

$$\sum_i E[\,{}_kR_i(t-\frac{\tau}{2})\ {}_kR_i(t+\frac{\tau}{2})] =$$

$$\sum_i E[{}_kR_i(t-\frac{\tau}{2})]\,E\,[{}_kR_i(t+\frac{\tau}{2})] - Cov[\,{}_kR_i(t-\frac{\tau}{2}),\,{}_kR_i(t+\frac{\tau}{2})] <$$

$$\sum_i E[{}_kR_i(t-\frac{\tau}{2})]\,E[{}_kR_i(t+\frac{\tau}{2})] =$$

$$\sum_i R_i\left(t-\frac{\tau}{2}\right)\,R_i\left(t+\frac{\tau}{2}\right) =$$

$$\hat{R}\left(t-\frac{\tau}{2}\right).\hat{R}(t+\frac{\tau}{2})$$

Therefore, $E\left[\widehat{{}_kR}\left(t-\frac{\tau}{2}\right).\widehat{{}_kR}\left(t+\frac{\tau}{2}\right)\right] < \hat{R}\left(t-\frac{\tau}{2}\right).\hat{R}(t+\frac{\tau}{2})$, i.e., the stability profile obtained using individual trial trajectories and then averaging across trials will be lower than that obtained from the expected single-trial trajectory (as we do here). The inequality holds only if the second term in the sum, the covariance between normalized time-separated activity across trials, is positive. Intuitively, this seems likely because of the fact that the activity of these neurons evolves smoothly. We verified this by computing the evolution of temporal covariance across trials for each neuron. Since we do not have access to the population unit vectors for individual trials, we simply used the unnormalized activity, reasoning that covariance should be unaffected by normalization scaling. The average covariance profile is shown in Figure A1, panel D. Note that the covariance is largely positive, with minor deviations late during the trial. Importantly, note that the covariance is higher during the visual epoch compared to the premotor epoch. This will produce a greater reduction in stability (step 4 in derivation) during the visual epoch in the true single-trial trajectories, only strengthening our result.

BIBLIOGRAPHY

Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV (2011) Single-trial neural correlates of arm movement preparation. Neuron 71:555-564.

Awh E, Armstrong KM, Moore T (2006) Visual and oculomotor selection: links, causes and implications for spatial attention. Trends Cogn Sci 10:124-130.

Balan PF, Ferrera VP (2003) Effects of gaze shifts on maintenance of spatial memory in macaque frontal eye field. J Neurosci 23:5446-5454.

Basso MA, Wurtz RH (1997) Modulation of neuronal activity by target uncertainty. Nature 389:66-69.

Berardelli A, Cruccu G, Manfredi M, Rothwell JC, Day BL, Marsden CD (1985) The corneal reflex and the R2 component of the blink reflex. Neurology 35:797-801.

Bergeron A, Guitton D (2000) Fixation neurons in the superior colliculus encode distance between current and desired gaze positions. Nature neuroscience 3:932-939.

Bergeron A, Matsuo S, Guitton D (2003) Superior colliculus encodes distance to target, not saccade amplitude, in multi-step gaze shifts. Nature neuroscience 6:404-413.

Borst JG, Sakmann B (1998) Facilitation of presynaptic calcium currents in the rat brainstem. The Journal of physiology 513 (Pt 1):149-155.

Bracewell RM, Mazzoni P, Barash S, Andersen RA (1996) Motor intention activity in the macaque's lateral intraparietal area. II. Changes of motor plan. J Neurophysiol 76:1457-1464.

Braff DL, Geyer MA (1990) Sensorimotor gating and schizophrenia - human and animal-model studies. Archives of General Psychiatry 47:181-188.

Branco T, Clark BA, Hausser M (2010) Dendritic discrimination of temporal input sequences in cortical neurons. Science 329:1671-1675.

Bryant CL, Gandhi NJ (2005) Real-time data acquisition and control system for the measurement of motor and neural data. J Neurosci Methods 142:193-200.

Buneo CA, Jarvis MR, Batista AP, Andersen RA (2002) Direct visuomotor transformations for reaching. Nature 416:632-636.

Buschman TJ, Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science 315:1860-1862.

Buschman TJ, Miller EK (2009) Serial, covert shifts of attention during visual search are reflected by the frontal eye fields and correlated with population oscillations. Neuron 63:386-396.

Card G, Dickinson MH (2008) Visually mediated motor planning in the escape response of Drosophila. Curr Biol 18:1300-1307.

Carello CD, Krauzlis RJ (2004) Manipulating intent: evidence for a causal role of the superior colliculus in target selection. Neuron 43:575-583.

Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements [see comments]. Nature 377:59-62.

Carter AG, Soler-Llavina GJ, Sabatini BL (2007) Timing and location of synaptic inputs determine modes of subthreshold integration in striatal medium spiny neurons. J Neurosci 27:8967-8977.

Castellote JM, Kumru H, Queralt A, Valls-Sole J (2007) A startle speeds up the execution of externally guided saccades. Experimental brain research 177:129-136.

Cavanaugh J, Wurtz RH (2004) Subcortical modulation of attention counters change blindness. J Neurosci 24:11236-11243.

Chang MH, Armstrong KM, Moore T (2012) Dissociation of response variability from firing rate effects in frontal eye field neurons during visual stimulation, working memory, and attention. J Neurosci 32:2204-2216.

Churchland MM, Santhanam G, Shenoy KV (2006a) Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. J Neurophysiol 96:3130-3146.

Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV (2006b) Neural variability in premotor cortex provides a signature of motor preparation. J Neurosci 26:3697-3712.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. Nature 487:51-56.

Churchland MM et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nature neuroscience 13:369-378.

Cisek P, Kalaska JF (2004) Neural correlates of mental rehearsal in dorsal premotor cortex. Nature 431:993-996.

Cohen B, Henn V (1972) Unit activity in the pontine reticular formation associated with eye movements. Brain Res 46:403-410.

Cohen JY, Pouget P, Heitz RP, Woodman GF, Schall JD (2009) Biophysical support for functionally distinct cell types in the frontal eye field. J Neurophysiol 101:912-916.

Colby CL, Goldberg ME (1999) Space and attention in parietal cortex. Annu Rev Neurosci 22:319-349.

Corbetta M, Akbudak E, Conturo TE, Snyder AZ, Ollinger JM, Drury HA, Linenweber MR, Petersen SE, Raichle ME, Van Essen DC, Shulman GL (1998) A common network of functional areas for attention and eye movements. Neuron 21:761-773.

Crawford TJ, Bennett D, Lekwuwa G, Shaunak S, Deakin JF (2002) Cognition and the inhibitory control of saccades in schizophrenia and Parkinson's disease. Prog Brain Res 140:449-466.

Curtis CE, D'Esposito M (2006) Selection and maintenance of saccade goals in the human frontal eye fields. J Neurophysiol 95:3923-3927.

Dorris MC, Munoz DP (1995) A neural correlate for the gap effect on saccadic reaction times in monkey. J Neurophysiol 73:2558-2562.

Dorris MC, Munoz DP (1998) Saccadic probability influences motor preparation signals and time to saccadic initiation. J Neurosci 18:7015-7026.

Dorris MC, Paré M, Munoz DP (1997) Neuronal activity in monkey superior colliculus related to the initiation of saccadic eye movements. J Neurosci 17:8566-8579.

Dorris MC, Olivier E, Munoz DP (2007) Competitive integration of visual and preparatory signals in the superior colliculus during saccadic programming. J Neurosci 27:5053-5062.

Edelman JA, Keller EL (1996) Activity of visuomotor burst neurons in the superior colliculus accompanying express saccades. J Neurophysiol 76:908-926.

Engbert R, Kliegl R (2003) Microsaccades uncover the orientation of covert attention. Vision Res 43:1035-1045.

Everling S, Paré M, Dorris MC, Munoz DP (1998) Comparison of the discharge characteristics of brain stem omnipause neurons and superior colliculus fixation neurons in monkey: implications for control of fixation and saccade behavior. J Neurophysiol 79:511-528.

Ferenczi E, Deisseroth K (2016) Illuminating next-generation brain therapies. Nature neuroscience.

Fischer B, Boch R (1983) Saccadic eye movements after extremely short reaction times in the monkey. Brain Res 260:21-26.

Fotowat H, Gabbiani F (2007) Relationship between the phases of sensory and motor activity during a looming-evoked multistage escape behavior. J Neurosci 27:10047-10059.

Funahashi S (2001) Neuronal mechanisms of executive control by the prefrontal cortex. Neurosci Res 39:147-165.

Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. Brain : a journal of neurology 119 (Pt 2):593-609.

Gandhi NJ, Keller EL (1999) Comparison of saccades perturbed by stimulation of the rostral superior colliculus, the caudal superior colliculus, and the omnipause neuron region. J Neurophysiol 82:3236-3253.

Gandhi NJ, Sparks DL (2004) Changing views of the role of the superior colliculus in the control of gaze. In: The Visual Neurosciences (Chalupa LM, Werner JS, eds), pp 1449-1465. Boston: MIT Press.

Gandhi NJ, Bonadonna DK (2005) Temporal interactions of air-puff evoked blinks and saccadic eye movements: Insights into motor preparation. J Neurophysiol 93:1718-1729.

Gandhi NJ, Katnani HA (2011) Motor functions of the superior colliculus. Annu Rev Neurosci 34:203-229.

Glimcher PW, Sparks DL (1992) Movement selection in advance of action in the superior colliculus. Nature 355:542-545.

Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. Nature 404:390-394.

Gold JI, Shadlen MN (2007) The neural basis of decision making. Annu Rev Neurosci 30:535-574.

Goldberg ME, Wurtz RH (1972) Activity of superior colliculus in behaving monkey. II. Effect of attention on neuronal responses. J Neurophysiol 35:560-574.

Goossens HH, Van Opstal AJ (2000a) Blink-perturbed saccades in monkey. I. Behavioral analysis. J Neurophysiol 83:3411-3429.

Goossens HH, Van Opstal AJ (2000b) Blink-perturbed saccades in monkey. II. Superior colliculus activity. J Neurophysiol 83:3430-3452.

Gregoriou GG, Gotts SJ, Desimone R (2012) Cell-type-specific synchronization of neural activity in FEF with V4 during attention. Neuron 73:581-594.

Hafed ZM, Goffart L, Krauzlis RJ (2008) Superior colliculus inactivation causes stable offsets in eye position during tracking. J Neurosci 28:8124-8137.

Hafed ZM, Goffart L, Krauzlis RJ (2009) A neural mechanism for microsaccade generation in the primate superior colliculus. Science 323:940-943.

Hafed ZM, Lovejoy LP, Krauzlis RJ (2013) Superior colliculus inactivation alters the relationship between covert visual attention and microsaccades. The European journal of neuroscience 37:1169-1181.

Hanes DP, Schall JD (1996) Neural control of voluntary movement initiation. Science 274:427-430.

Harvey MA, Saal HP, Dammann JF, 3rd, Bensmaia SJ (2013) Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex. PLoS biology 11:e1001558.

Hikosaka O, Nakamura K, Nakahara H (2006) Basal ganglia orient eyes to reward. J Neurophysiol 95:567-584.

Hikosaka O, Bromberg-Martin E, Hong S, Matsumoto M (2008) New insights on the subcortical representation of reward. Current opinion in neurobiology 18:203-208.

Hoffman JE, Subramaniam B (1995) The role of visual attention in saccadic eye movements. Percept Psychophys 57:787-795.

Horwitz GD, Newsome WT (1999) Separate signals for target selection and movement specification in the superior colliculus. Science 284:1158-1161.

Huerta MF, Harting JK (1984) The mammalian superior colliculus: Studies of its morphology and connections. In: Comparative Neurology of the Optic Tectum (Vanegas H, ed), pp 687-773. New York: Plenum Publishing.

Ignashchenkova A, Dicke PW, Haarmeier T, Thier P (2004) Neuron-specific contribution of the superior colliculus to overt and covert shifts of attention. Nature neuroscience 7:56-64.

Isa T, Hall WC (2009) Exploring the superior colliculus in vitro. J Neurophysiol 102:2581-2593.

Izawa Y, Suzuki H, Shinoda Y (2009) Response properties of fixation neurons and their location in the frontal eye field in the monkey. J Neurophysiol 102:2410-2422.

Jagadisan UK, Gandhi NJ (2016) Disruption of Fixation Reveals Latent Sensorimotor Processes in the Superior Colliculus. J Neurosci 36:6129-6140.

Jantz JJ, Watanabe M, Everling S, Munoz DP (2013) Threshold mechanism for saccade initiation in the frontal eye field and superior colliculus. J Neurophysiol 109:2767-2780.

Juan CH, Shorter-Jacobi SM, Schall JD (2004) Dissociation of spatial attention and saccade preparation. Proceedings of the National Academy of Sciences of the United States of America 101:15541-15544.

Katnani HA, Gandhi NJ (2013) Time course of motor preparation during visual search with flexible stimulus-response association. J Neurosci 33:10057-10065.

Katnani HA, Van Opstal AJ, Gandhi NJ (2012) Blink perturbation effects on saccades evoked by microstimulation of the superior colliculus. PLoS One 7:e51843.

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. Nature neuroscience 17:440-448.

Keller EL (1974) Participation of medial pontine reticular formation in eye movement generation in monkey. J Neurophysiol 37:316-332.

Kosslyn SM, Ganis G, Thompson WL (2001) Neural foundations of imagery. Nat Rev Neurosci 2:635-642.

Koval MJ, Lomber SG, Everling S (2011) Prefrontal cortex deactivation in macaques alters activity in the superior colliculus and impairs voluntary control of saccades. J Neurosci 31:8659-8668.

Kraskov A, Dancause N, Quallo MM, Shepherd S, Lemon RN (2009) Corticospinal neurons in macaque ventral premotor cortex with mirror properties: a potential mechanism for action suppression? Neuron 64:922-930.

Kraskov A, Philipp R, Waldert S, Vigneswaran G, Quallo MM, Lemon RN (2014) Corticospinal mirror neurons. Philosophical transactions of the Royal Society of London Series B, Biological sciences 369:20130174.

Krauzlis RJ, Liston D, Carello CD (2004) Target selection and the superior colliculus: goals, choices and hypotheses. Vision Res 44:1445-1451.

Kumru H, Valls-Sole J (2006) Excitability of the pathways mediating the startle reaction before execution of a volutary movement. Experimental brain research 169:427-432.

Kustov AA, Robinson DL (1996) Shared neural control of attentional shifts and eye movements. Nature 384:74-77.

Lo CC, Wang XJ (2006) Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. Nature neuroscience 9:956-963.

Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503:78-84.

May PJ (2006) The mammalian superior colliculus: laminar structure and connections. Progress in brain research 151:321-378.

Mazzoni P, Bracewell RM, Barash S, Andersen RA (1996) Motor intention activity in the macaque's lateral intraparietal area. I. Dissociation of motor plan from sensory memory. J Neurophysiol 76:1439-1456.

McPeek RM, Keller EL (2002) Saccade target selection in the superior colliculus during a visual search task. J Neurophysiol 88:2019-2034.

Meredith MA, Ramoa AS (1998) Intrinsic circuitry of the superior colliculus: pharmacophysiological identification of horizontally oriented inhibitory interneurons. J Neurophysiol 79:1597-1602.

Moore T, Fallah M (2001) Control of eye movements and spatial attention. Proceedings of the National Academy of Sciences of the United States of America 98:1273-1276.

Moore T, Armstrong KM (2003) Selective gating of visual signals by microstimulation of frontal cortex. Nature 421:370-373.

Moschovakis AK, Scudder CA, Highstein SM (1996) The microscopic anatomy and physiology of the mammalian saccadic system. Prog Neurobiol 50:133-254.

Muller JR, Philiastides MG, Newsome WT (2005) Microstimulation of the superior colliculus focuses attention without moving the eyes. Proceedings of the National Academy of Sciences of the United States of America 102:524-529.

Munoz DP, Wurtz RH (1993a) Fixation cells in monkey superior colliculus. I. Characteristics of cell discharge. J Neurophysiol 70:559-575.

Munoz DP, Wurtz RH (1993b) Fixation cells in monkey superior colliculus. II. Reversible activation and deactivation. J Neurophysiol 70:576-589.

Munoz DP, Istvan PJ (1998) Lateral inhibitory interactions in the intermediate layers of the monkey superior colliculus. J Neurophysiol 79:1193-1209.

Munoz DP, Waitzman DM, Wurtz RH (1996) Activity of neurons in monkey superior colliculus during interrupted saccades. J Neurophysiol 75:2562-2580.

Munoz DP, Armstrong IT, Hampton KA, Moore KD (2003) Altered control of visual fixation and saccadic eye movements in attention-deficit hyperactivity disorder. J Neurophysiol 90:503-514.

Nakagawa H, Nishida Y (2012) Motor planning modulates sensory-motor control of collision avoidance behavior in the bullfrog, Rana catesbeiana. Biol Open 1:1094-1101.

Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. Nature 341:52-54.

Ottes FP, Van Gisbergen JA, Eggermont JJ (1986) Visuomotor fields of the superior colliculus: a quantitative model. Vision Res 26:857-873.

Phongphanphanee P, Marino RA, Kaneda K, Yanagawa Y, Munoz DP, Isa T (2014) Distinct local circuit properties of the superficial and intermediate layers of the rodent superior colliculus. The European journal of neuroscience 40:2329-2343.

Pierrot-Deseilligny C, Rivaud S, Gaymard B, Agid Y (1991) Cortical control of reflexive visually-guided saccades. Brain : a journal of neurology 114 (Pt 3):1473-1485.

Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. Nature 400:233-238.

Preuss T, Osei-Bonsu PE, Weiss SA, Wang C, Faber DS (2006) Neural representation of object approach in a decision-making motor circuit. J Neurosci 26:3454-3464.

Ratcliff R, Rouder JN (1998) Modeling response times for two-choice decisions. Psychol Rev 9:347-356.

Ratcliff R, Cherian A, Segraves M (2003) A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. J Neurophysiol 90:1392-1407.

Rauch A, La Camera G, Luscher HR, Senn W, Fusi S (2003) Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. J Neurophysiol 90:1598-1612.

Reddi BA, Asrress KN, Carpenter RH (2003) Accuracy, information, and response time in a saccadic decision task. J Neurophysiol 90:3538-3546.

Reuter B, Kathmann N (2004) Using saccade tasks as a tool to analyze executive dysfunctions in schizophrenia. Acta Psychol (Amst) 115:255-269.

Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. Nature 497:585-590.

Rizzolatti G, Riggio L, Dascola I, Umilta C (1987) Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. Neuropsychologia 25:31-40.

Robinson DA (1972) Eye movements evoked by collicular stimulation in the alert monkey. Vision Res 12:1795-1808.

Rodgers CK, Munoz DP, Scott SH, Pare M (2006) Discharge properties of monkey tectoreticular neurons. J Neurophysiol 95:3502-3511.

Rottach KG, Das VE, Wohlgemuth W, Zivotofsky AZ, Leigh RJ (1998) Properties of horizontal saccades accompanied by blinks. J Neurophysiol 79:2895-2902.

Sadeh M, Sajad A, Wang H, Yan X, Crawford JD (2015) Spatial transformations between superior colliculus visual and motor response fields during head-unrestrained gaze shifts. The European journal of neuroscience 42:2934-2951.

Sajad A, Sadeh M, Keith GP, Yan X, Wang H, Crawford JD (2015) Visual-Motor Transformations Within Frontal Eye Fields During Head-Unrestrained Gaze Shifts in the Monkey. Cerebral cortex 25:3932-3952.

Sato TR, Schall JD (2003) Effects of stimulus-response compatibility on neural selection in frontal eye field. Neuron 38:637-648.

Schall JD (2004) On the role of frontal eye field in guiding attention and saccades. Vision research 44:1453-1467.

Schall JD, Hanes DP (1993) Neural basis of saccade target selection in frontal eye field during visual search. Nature 366:467-469.

Schall JD, Purcell BA, Heitz RP, Logan GD, Palmeri TJ (2011) Neural mechanisms of saccade target selection: gated accumulator model of the visual-motor cascade. Eur J Neurosci 33:1991-2002.

Schultz KP, Williams CR, Busettini C (2010) Macaque pontine omnipause neurons play no direct role in the generation of eye blinks. J Neurophysiol 103:2255-2274.

Segraves MA (1992) Activity of monkey frontal eye field neurons projecting to oculomotor regions of the pons. J Neurophysiol 68:1967-1985.

Shen K, Pare M (2007) Neuronal activity in superior colliculus signals both stimulus identity and saccade goals during visual conjunction search. J Vis 7:15 11-13.

Softky W (1994) Sub-millisecond coincidence detection in active dendritic trees. Neuroscience 58:13-41.

Sommer MA, Wurtz RH (2001) Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. J Neurophysiol 85:1673-1685.

Sparks DL, Hartwich-Young R (1989) The deep layers of the superior colliculus. Rev Oculomot Res 3:213-255.

Sparks DL, Mays LE (1990) Signal transformations required for the generation of saccadic eye movements. Annu Rev Neurosci 13:309-336.

Sparks DL, Holland R, Guthrie BL (1976) Size and distribution of movement fields in the monkey superior colliculus. Brain Res 113:21-34.

Steinmetz NA, Moore T (2010) Changes in the response rate and response variability of area V4 neurons during the preparation of saccadic eye movements. J Neurophysiol 103:1171-1178.

Steinmetz NA, Moore T (2012) Lumping and splitting the neural circuitry of visual attention. Neuron 73:410-412.

Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J (2013) Dynamic coding for cognitive control in prefrontal cortex. Neuron 78:364-375.

Terao Y, Fukuda H, Yugeta A, Hikosaka O, Nomura Y, Segawa M, Hanajima R, Tsuji S, Ugawa Y (2011) Initiation and inhibitory control of saccades with the progression of Parkinson's disease - Changes in three major drives converging on the superior colliculus. Neuropsychologia 49:1794-1806.

Thompson KG, Biscoe KL, Sato TR (2005) Neuronal basis of covert spatial attention in the frontal eye field. J Neurosci 25:9479-9487.

Thompson KG, Hanes DP, Bichot NP, Schall JD (1996) Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. J Neurophysiol 76:4040-4055.

Tsujimoto T, Jeromin A, Saitoh N, Roder JC, Takahashi T (2002) Neuronal calcium sensor 1 and activity-dependent facilitation of P/Q-type calcium currents at presynaptic nerve terminals. Science 295:2276-2279.

Vigneswaran G, Philipp R, Lemon RN, Kraskov A (2013) M1 corticospinal mirror neurons and their role in movement suppression during action observation. Curr Biol 23:236-243.

Walton MM, Gandhi NJ (2006) Behavioral evaluation of movement cancellation. J Neurophysiol 96:2011-2024.

Winterson BJ, Collewijn H (1976) Microsaccades during finely guided visuomotor tasks. Vision research 16:1387-1390.

Wurtz RH, Sommer MA, Paré M, Ferraina S (2001) Signal transformations from cerebral cortex to superior colliculus for the generation of saccades. Vision Res 41:3399-3412.

Zandbelt B, Purcell BA, Palmeri TJ, Logan GD, Schall JD (2014) Response times from ensembles of accumulators. Proceedings of the National Academy of Sciences of the United States of America 111:2848-2853.

Zhang M, Barash S (2004) Persistent LIP activity in memory antisaccades: working memory for a sensorimotor transformation. J Neurophysiol 91:1424-1441.

Zuber BL, Stark L (1965) Microsaccades and velocity-amplitude relationship for saccadic eye movements. Science 150:1459-1460.