Neural Population Dynamics of Sensorimotor Signals for Eye Movements

by

Michelle Renae Heusser

B.S. in Biomedical Engineering, University of Arizona, 2015

Submitted to the Graduate Faculty of the

Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2022

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Michelle Renae Heusser

It was defended on

June 3, 2022

and approved by

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

Matthew Smith, Ph.D., Associate Professor, Department of Biomedical Engineering and
Neuroscience Institute, Carnegie Mellon University

Byron Yu, Ph.D., Professor, Departments of Electrical and Computer Engineering and
Biomedical Engineering, Carnegie Mellon University

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

Neural Population Dynamics of Sensorimotor Signals for Eye Movements

Michelle R. Heusser, PhD

University of Pittsburgh, 2022

During active vision, we convert information about visual objects in our periphery into goal-directed eye movements known as saccades. This process of sensorimotor integration is complex; we must incorporate knowledge about our environment, including the spatial location of the target object and the urgency of saccade initiation. The superior colliculus (SC) is a deep brain structure that is critical for active vision, with most neurons in this area responding to the presence of a visual stimulus and increasing their activity to signal for saccade initiation. In the studies presented in this dissertation, we characterized the combined activity patterns of small populations of neurons in the non-human primate SC across multiple contexts to probe various parameters of active vision. We used simple machine learning techniques (i.e., dimensionality reduction and/or classification) that quantitatively capture the activity pattern across many simultaneously recorded channels. First, we examined the dynamics of population activity during the time between sensation and action and found that activity slowly evolves from a visual-like to a motor-like pattern when a delay is imposed. This sensorimotor transformation signature is robust to perturbations induced by small fixational saccades and is correlated with saccade latency, indicative of a potential mechanism for movement generation. Next, we investigated the impact of behavioral context on the population-level representation during the sensation and action periods of active vision and observed unique encoding of both content (sensation/action epochs) and context (two comparable behavioral tasks). Last, we determined the time course and spatial extent of intended saccade target direction encoding by SC neural populations in an eight-target delayed saccade task. We compared these profiles with a second signal modality – the local field potentials (LFPs), which represent collective activity in a broader region of the SC. Neural spiking activity better encoded target direction throughout the time course of sensorimotor integration than did LFP signals. Population activity during the motor epoch exhibited broader spatial tuning than in the visual epoch, indicative of dynamic encoding of spatial parameters. Taken together, these studies provide foundational knowledge of the SC’s role in the process of active vision.

Table of Contents

List of Figures

Acknowledgements

I have received an abundance of support throughout my graduate career, so here I will make an attempt to express my immense gratitude. First, I'd especially like to thank my advisor, Dr. Neeraj (Raj) Gandhi, for his mentorship. He has a wealth of knowledge and experience that he freely shares, and I am grateful that I have learned so much from him during my time working in his lab. Raj truly cares about his trainees and will actively go out of his way to provide whatever type of support he can, whether it be academic, professional, or personal in nature.

I'd also like to thank my committee members – Drs. Aaron Batista, Matthew Smith, and Byron Yu – for their valuable discussions and guidance over the years. Also, the ARCS Foundation provided me more than just a scholarship; I became part of a wonderful community in Pittsburgh and have met so many women who value supporting young scientists like myself. The donors of my award – Kate Freed, Sarah Jones, and Dr. Pamela Meadowcroft – are my cheerleaders; I cherish their enthusiasm and kind hearts.

I am thankful for the support of my labmates, and I'd like to shout out a few. Dr. Clara Bourrelly collected much of the data included in this dissertation. Dr. Uday Jagadisan provided mentorship during the first few years of my program. Dr. Ivan Smalianchuk never hesitated to help me troubleshoot MATLAB code or discuss random science topics with me and Uday while chaotically scribbling on every last inch of space on our lab's whiteboard to brainstorm. Eve Ayar has been my undergraduate mentee for the past two or so years. She is dedicated, taking and running with our seedlings of ideas to turn them into full-fledged projects.

In addition to thanking my labmates for being great colleagues, many of them have become good friends over the years. This lab has been a close-knit group, and hanging out together outside of the lab environment has made my grad student experience that much better. I fondly remember going on runs and working out with Uday, Kevin, and Clara – especially when we would bombard Corentin's work area and turn it into a "Gandhi Lab Locker Room." Having these workout buddies motivated me to run my first 10K with a personal-best pace. I always enjoy hanging out and playing video or board games with Ivan (and also have to thank him for introducing me to my partner!). I value my trips to the Indian food truck with Eve and the chats we have along the way.

Next, I'd like to thank all the friends I've made during my time here – those in student groups like BMES and EGSO, and those in other parts of the program. I'd especially like to acknowledge the support I've had in my closest friends, Dr. Erinn Grigsby, Dr. Sally Zheng, and (soon-to-be Dr.) Nicole McClain, who all entered the BioE PhD program along with me. I'd say we've grown together and bonded over chats, hangouts, and knitting. I value their friendship so much. I also want to thank Rachel Blaine, who is an endless source of encouragement and motivation. She and I have been friends since our middle school youth orchestra days, and I'm so glad we've not only stayed in touch, but stayed close.

Next, I'd like to acknowledge my family, including my mom (Kathy), my dad (Ken), and my sister (Rachel). I can say confidently that I wouldn't be where I am today without my parents, as they've sacrificed so much to make sure I am set up for success in anything I do. Their dedication is unmatched. And of course, my sister is my lifelong best bud. We are able to be silly and goofy together, and we have always had each other's back.

Lastly, I'd like to thank my partner of almost 6 years, Dr. Corey Karadeema. He has seen both the ups and the downs of grad school and life, and I love all the adventures we've shared. He's the most loyal, caring, and patient person I've ever met, and I thank him for everything he is and everything he's done for me.

1.0 Introduction

Active vision is a complex behavioral phenomenon. When we see an object in the periphery of our visual field, we often make rapid eye movements known as saccades to redirect our line of sight to the visual stimulus to project its image on the fovea where visual acuity is the highest. The process of converting visual information into a goal-directed eye movement is also called *sensorimotor transformation*. The intricacies of this process in natural settings are perhaps obvious – we must take into account the distance and direction of the visual target from our current eye position, incorporate our prior understanding or estimation of the object's identity, and often withhold or delay an eye movement when appropriate. The studies presented in this dissertation isolate three different aspects of active vision and investigate the manners by which groups of neurons in one brain area (henceforth referred to as *neural populations*) represent their activity during the process of sensorimotor transformation.

1.1 The Superior Colliculus, a Key Player in Active Vision

The superior colliculus (SC) is a neural structure in the midbrain that is integral to the process of converting visual information into an eye movement command. Located close to the effectors (i.e., eye muscles), this area is late in the oculomotor circuitry and receives signals from almost every structure involved in eye movements (Wurtz et al., 2001). Many SC neurons transiently increase their activity level before the onset of a saccade ("motor burst"). Some neurons fire in a similar manner following the appearance of a visual stimulus in the periphery ("visual burst"), with the majority of neurons signaling to some degree during both time periods, or *epochs* (Wurtz et al., 2001). Visual, visuomotor, and motor neurons vary along the dorsoventral axis of the SC in a continuum, with visual neurons found in the superficial layers, visuomotor neurons in the intermediate and deep layers, and motor neurons in the deep layers (Basso & May, 2017; Massot et al., 2019).

Within a single layer (i.e., the plane orthogonal to the dorsoventral axis), neurons in the SC are topographically organized – the spatial tuning preferences of neurons are indicated by their position along the rostral-caudal and mediolateral axes. Along these axes, neurons that are visually responsive fire maximally when a target is presented in the preferred location (i.e., target eccentricity and direction, respectively, relative to current eye position), and neurons that are motor-related fire similarly when signaling to generate a saccade with their preferred spatial parameters (again, to their preferred amplitude and direction, respectively) (Gandhi & Katnani, 2011). Figure 1 shows this spatial topography. In addition to the strong relationship observed between SC activity and the external variables of sensation or action, the SC has also been implicated in the processing of cognitive factors such as attention (Lovejoy & Krauzlis, 2009) and decision-making (Crapse et al., 2018; Keller et al., 2005). An important advance of the studies presented in this dissertation is our population-level approach to linking SC activity to several behavioral parameters, as described next.

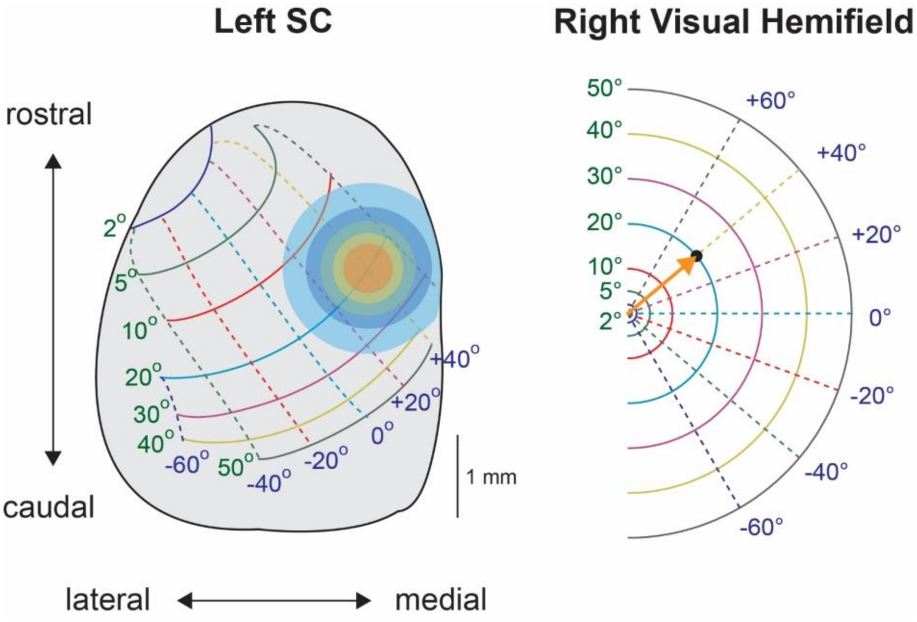

Figure 1. Depiction of SC topography for spatial parameters.

Left: One layer of the SC and the firing properties of the neurons at different locations on the SC map. Along the rostral-caudal axis, neurons vary systematically in their preferred visual target eccentricity, and equivalently, in their preferred saccade amplitude. Along the mediolateral axis, neurons vary systematically in their preferred visual stimulus direction (measured in degrees of visual angle), and likewise for their preferred direction of the executed saccade. Right: The hemifield of visual space to which the left SC is related. For a given visual stimulus location or saccade vector (indicated by the black dot and the orange arrow, respectively), neurons located at the equivalent region on the SC map are maximally active, with surrounding neurons firing sub-optimally, leading to a "Gaussian mound" of activation (represented here by concentric rings on the SC map). Figure adapted from Gandhi & Katnani, 2011.

1.2 Identifying Neural Correlates of Behavioral Phenomena

Historically, our knowledge of the link between brain areas and behavioral parameters has been limited to results obtained from single-unit studies. In these studies, experimenters acutely insert a single electrode, find and isolate an individual neuron, and record its spiking activity while the animal performs many repetitions ("trials") of a behavioral task. This process is repeated for tens to hundreds of days to obtain a neural population. The activity is then temporally aligned to some event of interest (e.g., visual stimulus appearance) and averaged across many repetitions of the behavioral task and across all neurons in the population to identify modulations of activity that relate to the parameter of interest (e.g., attention).

Recently, a dynamical systems approach has taken the neuroscience field by storm. With advances in technology, experimenters can record the activity of hundreds (for cortical structures) or tens (for deep brain structures) of neurons simultaneously with implanted or acutely inserted laminar electrode arrays, respectively. We use this advance to our advantage; instead of averaging activity across trials and neurons, we retain and leverage information about the correlational structure across all simultaneously recorded neurons that produces a single measure of their combined activity patterns *throughout single trials* (Shenoy et al., 2013). The following subsection describes our methodology for this holistic population-level approach to investigating sensorimotor encoding properties in the SC.

1.3 Machine Learning for Analysis of Population Dynamics

1.3.1 Dimensionality Reduction

Consider the following hypothetical case – you record the activity of two neurons before and after a visual stimulus is flashed in an animal's periphery. Neuron A increases its activity by 50 spikes per second (spk/s) while the activity of neuron B is suppressed by 50 spk/s compared to its activity rate at rest, or *baseline*. If the two neurons modulate their activity in (roughly equal and) opposite ways in response to this behavioral parameter, the average activity modulation across your population is 0 spk/s. You would not be able to link the combined activity of these neurons to a behavioral feature because this averaging measure washes out any effect. If instead, the modulation of the two neurons considered above was represented as a vector [50 -50], all information about the direction (elevation or suppression of activity) as well as the magnitude (amount of activity change in units of spk/s) of the response is retained. Extrapolating to a much larger population of *N* simultaneously recorded neurons, we can think of the population response as a vector in an *N*-dimensional space. The pattern of responses across all neurons is much more likely to be informative about the neural responses underlying motor control than is an averaging measure. But how can we represent this population activity pattern in a simple and intuitive way? It is hard to think in *N*-dimensional spaces. However, it is easy to picture that there are fewer degrees of freedom in the population response than there are neurons if we realize that these neurons are all connected in a network. For example, if the animal repeats the same eye movement a few times, the neural population response won't be exactly [50, -50] each time; there will be some variability. But again, because the neurons in question are part of the same network, their variability will shared – that is, not independent. This is where dimensionality reduction comes in. When applied to neural data, dimensionality reduction algorithms provide a summary of population activity by finding weighted linear combinations of neurons that relate to the largest amount of shared variability in neural activity across trials.

In recent years, the field has harnessed these intuitions to capture low-dimensional structure in neural population responses – responses that often exhibit temporal dynamics, just as in many other physical or engineered systems. Many research groups will first reduce the dimensionality of their neural data through one of many possible algorithms (e.g., principal components analysis or factor analysis) and study the activity pattern in the resulting low-dimensional “state space” (see Figure 2). The temporal evolution of the activity position is indicative of the activity pattern modulation surrounding behavioral events of interest. In Chapters 2 and 3, we applied this concept to study the patterns of population activity in the oculomotor system during a sensorimotor behavioral task. More specifically, we utilized an algorithm called Gaussian Process Factor Analysis (GPFA), which concurrently optimizes the low-dimensional state space parameters and the amount of smoothing appropriate for the input activity (Yu et al., 2009).

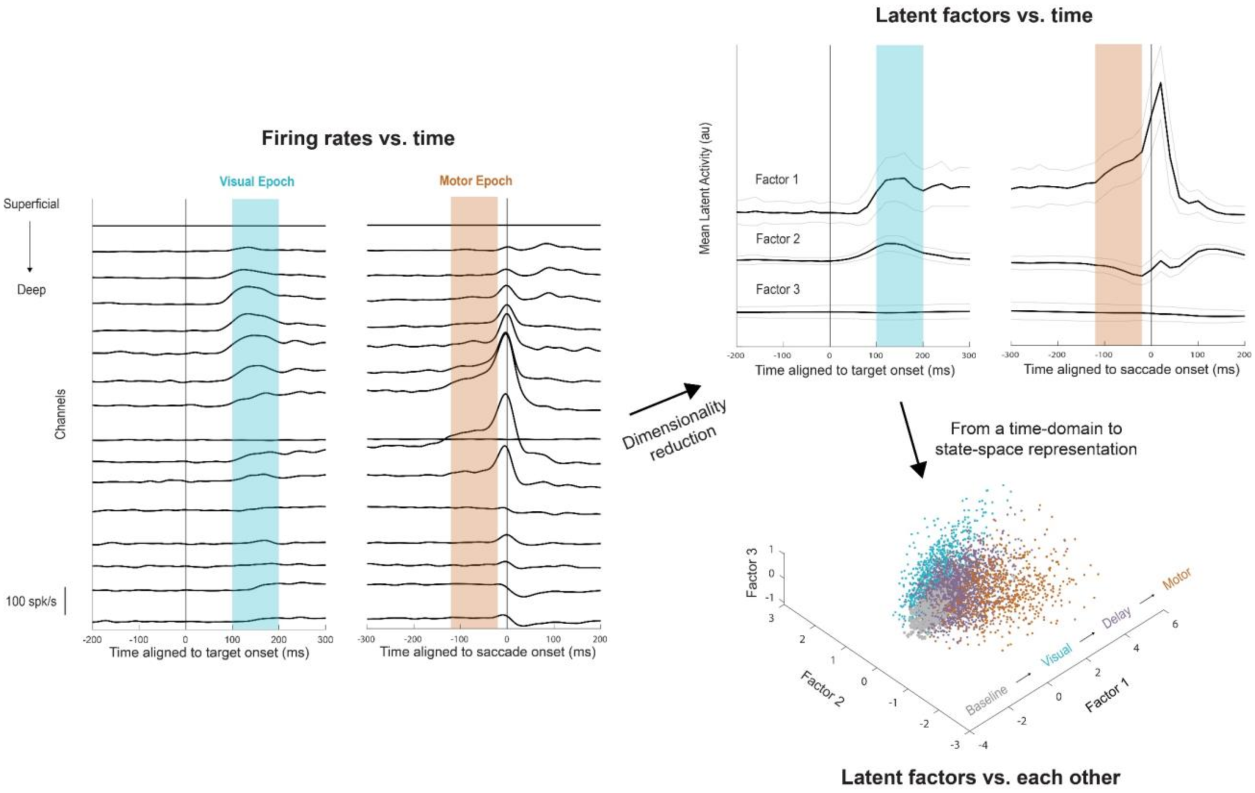

Figure 2. Schematic of dimensionality reduction for population-level analysis of SC activity. Left: Representative neural activity profiles across SC depth (from the superficial layers at top to the deeper layers on bottom) aligned to the onset of a visual stimulus and to saccade onset during a simple behavioral task. Top right: After dimensionality reduction, a smaller number of latent factors have been identified. Each factor represents a linear summary of neural activity across all recorded channels. Bottom right: Latent activity in three dimensions plotted against each other rather than against time. This state-space representation allows for visualization and analysis of the activity pattern across the entire population at any given time. Figure adapted from Figure 4 in Chapter 2, and additional details can be found there.

1.3.2 Offline Decoding

Decoding is another machine learning technique useful for characterization of population activity patterns. Decoding was first introduced in the context of neural prosthetics, where the desire is to infer, or "decode," the movement intentions of a paralyzed person from their neural activity alone. Offline decoding algorithms, one version of which are *classifiers*, are often used as a tool to quantify the amount of information encoded by neural activity in relation to external parameters. In other words, the more distinct the neural activity patterns are across two or more conditions (for example, looking rightward vs. looking leftward), the more accurate the classifier will be at categorizing novel activity patterns as correctly belonging to a given condition, and thus, the more information (e.g., about target direction) is assumed to be encoded by the neurons. In Chapters 2 and 3, we implement a simple linear discriminant analysis (LDA) classifier to quantify the uniqueness of dimensionality-reduced population activity patterns in the SC between two epochs (sensation and action) and between two behavioral tasks (the delayed saccade task and the gap task). In Chapter 4 we extend the classification of population activity patterns to eight categories, each representing activity evoked for trials to a particular saccade target direction. From this simple classification method, we characterized the time course of spatial information encoding in SC signals throughout a sensorimotor task.

1.4 Dissertation Structure

In this dissertation, I examine how small populations of neurons in the primate SC encode sensorimotor, contextual, and spatial parameters during behavioral tasks in which visual information is converted into a goal-directed eye movement. Chapter 2 characterizes the time course of sensorimotor transformation signals present in SC neural populations and their relationship to movement preparation. Chapter 3 addresses the ways in which neural populations vary their activity patterns under two behavioral contexts. Chapter 4 challenges our existing notion that SC neurons in isolated areas on the SC topographic map lack information about spatial parameters across a large range of the visual field. The dissertation concludes in Chapter 5 with a summary of our findings and a look into future lines of research that could stem from these studies.

2.0 Drifting Representation with Transient Resets Characterizes Sensorimotor Transformation in the Monkey Superior Colliculus

This work was submitted as a manuscript: **Heusser MR**, Jagadisan UK, Gandhi, NJ. “Drifting representation with transient resets characterizes sensorimotor transformation in the monkey superior colliculus” (In Review)

2.1 Overview

To produce goal-directed eye movements known as saccades, we must channel sensory input from our environment through a process known as sensorimotor transformation. The behavioral output of this phenomenon (an accurate eye movement) is straightforward, but the coordinated activity of neurons underlying it is not well understood. We searched for a neural correlate of sensorimotor transformation in the activity patterns of simultaneously recorded neurons in the superior colliculus (SC) of rhesus monkeys performing a standard delayed saccade task. Neurons in its intermediate layers produce a burst of spikes both following the appearance of a visual (sensory) stimulus and preceding an eye movement command, but many also exhibit a sustained activity level during the intervening time (“delay period”). Each session’s population activity was summarized in a low-dimensional framework and assessed on a scale of visual- to motor-like throughout the delay period using a novel measure we call the Visuomotor Proximity Index (VMPI). On average, population activity slowly evolved from a more visual- to a more motor-like pattern throughout the delay period, but microsaccade perturbations transiently deviated it to a visual-like pattern. A correlation was also found between the VMPI and single trial saccadic reaction time, even hundreds of milliseconds before the cue to initiate a movement. Therefore, we conclude that SC population activity contains a neural signature of the sensorimotor transformation process, systematically drifting toward a motor-like representation and intermittently reverting to a visual-like representation following a microsaccade.

2.2 Introduction

Sensorimotor transformation is the framework by which our brains process sensory input and subsequently produce a motor command. Its functionality is easily appreciated in the oculomotor system – when we see an object in our periphery, we can promptly direct our line of sight to that target. However, at what times are the neural populations representing the presence of a visual target through their coordinated activity? At what times are they collectively producing a signal that more closely resembles a motor command? And how does the population response transition from sensory to motor representations?

The superior colliculus (SC) is a midbrain structure crucial for sensorimotor transformation (Basso and May, 2017; Cooper and McPeek, 2021; Gandhi and Katnani, 2011; Sajad et al., 2020; Wurtz and Optican, 1994). Neurons in its deeper layers emit strong bursts of activity both when a visual stimulus appears as well as when a high-velocity eye movement, known as a saccade, is generated to redirect gaze toward that object of interest. These putative “visual” and “motor” bursts are well characterized but the time course of integrating visual stimulus-related information into a motor command is not understood as well. Previous research on sensorimotor integration in the oculomotor system has relied on single unit studies. These studies have focused on reference frame transformations, with the objective of determining whether the temporally evolving neural activity better represents stimulus location or movement amplitude. The general result is that, immediately after stimulus presentation, the sensory response is encoded in the reference frame of the stimulus modality – oculocentric for vision and craniocentric for audition. Just prior to the movement onset, the activity is best represented as a motor command in eye-centered coordinates or in a hybrid reference frame. In the intervening delay period, the average activity shows a slow and systematic transition from sensation to action representations, one which is sped up when no delay period is imposed. Such findings have been reported in the SC (Lee and Groh, 2012; Sajad et al., 2020; Sadeh et al., 2020), frontal eye fields (Caruso et al., 2018b; Sajad et al., 2016), parietal cortex (Buneo et al., 2002; Mullette-Gillman et al., 2005), and supplementary eye fields (Bharmauria et al., 2021). We sought to characterize at a population level the moment-by-moment representation of SC activity between sensation and action. We labeled the transient burst representations that follow target onset and precede saccade onset as 'visual' and 'motor' subspaces, respectively, while remaining agnostic to their preferred coordinate system. We then sought to determine how the population activity during the delay period transitioned between the two representations. This approach provides a more unsupervised yet still direct understanding of how SC populations encode these features. To this end, we searched for a neural correlate of sensorimotor transformation in small populations of SC neurons by characterizing the “visual-like” or “motor-like” pattern of activity during the intervening period of time between the visual and motor bursts while rhesus monkeys (*Macaca mulatta*) performed a visually-guided delayed saccade task (Figure 3). This paradigm temporally separates the visual from the motor epoch through a “delay period” and has been previously employed in countless studies of cognition, sensation, and motor behavior.

To characterize the shared activity patterns of neural populations, machine learning methods such as dimensionality reduction have been utilized to investigate the dynamics of neural activity underlying cognitive or behavioral processes such as stimulus encoding (e.g., Cowley et al., 2016), decision making (e.g., Aoi et al., 2020), and movement execution (e.g., Churchland et al., 2006). Such techniques transform the activity across the population into a state-space framework, where the pattern at any given moment can be represented as a linear combination of the activity of individual neurons. This methodology offers a noise-reduced, better-visualizable trajectory of activity across consecutive time points (Cunningham and Yu, 2014). Here, we employed a dimensionality reduction algorithm called Gaussian Process Factor Analysis (GPFA, Yu et al., 2009), to characterize the time course of population-level representations as they relate to vision and saccadic eye movement. First, we used a linear discriminant analysis (LDA) classifier to determine if the "subspaces" formed by collective activity patterns during the visual and motor epochs were distinguishable from each other and found that for the bulk of neural populations, this was indeed the case. Exploiting this separability, we then computed the similarity of the activity patterns throughout the delay period to either the visual or the motor subspace through a Visuomotor Proximity Index (VMPI) (based on the proximity index in Dekleva et al., 2018). When looking across repetitions of the task, activity patterns exhibited a slow, systematic drift from a visual- to a motor-like pattern. Remarkably, whenever a microsaccade occurred during the delay period, the population activity pattern transiently deviated to a visual-like representation before rapidly returning to the original trajectory. Finally, we tested an existing theory of arm movement generation known as the "initial condition hypothesis" (Afshar et al., 2011) and found that the state-space position of the activity on a given trial was correlated with the eventual saccadic reaction time, a relationship that emerged even hundreds of milliseconds before the cue to initiate a movement.

This study extends our knowledge of the SC’s role in sensorimotor transformation through both a network-level analysis of neural activity across sensorimotor epochs as well as a direct investigation of the relationship between this intermediate activity and behavior. Taken all together, these findings indicate that 1) there is a neural signature of the sensorimotor transformation process present in SC populations that can be characterized by a slow drift with transient resets, and that 2) activity patterns that drift to a stronger motor-like representation by the end of the delay period may enable a more rapid initiation of a saccade, substantiating the idea that this movement initiation mechanism is conserved across motor systems.

2.3 Methods

2.3.1 Subjects and Surgical Approach

Three adult male rhesus monkeys (*Macaca mulatta;* monkeys BL, BB, and SU) were used for this study. The experimental protocol was approved by the University of Pittsburgh Institutional Animal Care and Use Committee. Each animal underwent a sterile surgery under general anesthesia to implant a cylindrical recording chamber (Narishige) positioned above a craniotomy that allows access to the SC. A Teflon-coated, stainless-steel wire was also implanted on one eye in some animals. Surgical methods are described in more detail in Jagadisan & Gandhi, 2016.

2.3.2 Visual Stimuli and Behavioral Paradigm

Stimulus presentation and the animal's behavior were under real-time control with a LabVIEW-based controller interface (Bryant & Gandhi, 2005). All stimuli were white squares, 4x4 pixels subtending approximately 0.5°, displayed against a dark grey background on a LED-backlit flat screen monitor. Eye position was recorded using the scleral search coil technique (CNC Engineering) or using an EyeLink 1000 eye tracker (SR Research), both sampled at 1 kHz.

Each monkey was trained to sit head-restrained in a primate chair and perform a standard delayed saccade task in a dimly lit room. To complete a successful trial of this task, the monkey fixated on a visual stimulus located in the center of the screen and maintained fixation while a visual target was presented in the animal's periphery. After a variable delay period (600-1200 ms for monkeys BL and BB, and 700-1500 ms for monkey SU, weighted to maintain a flat anticipation function), the fixation point was extinguished, serving as the animal's "go cue" to make a saccadic eye movement to the target. The animal had to make a saccade to the peripheral target within 460-800 ms and was required to maintain fixation on it for 300 ms to receive a liquid reward. The monkeys performed this task with high accuracy before recording sessions began. Thus, we limited our analyses to rewarded trials only.

2.3.3 Electrophysiology and Data Pre-processing

On each recording session, a 16- or 24-channel linear microelectrode array (Plexon or AlphaOmega) was inserted orthogonal to the SC surface along the dorsoventral axis. Neurons recorded using this approach had similar preferred saccade vectors as determined by microstimulation (Massot et al., 2019). Care was taken to position the electrode in a way that maximized the yield of neurons exhibiting both visual and motor bursts, and thereby tended to be positioned in the deeper SC layers.

On each given trial, the target presented in the animal's periphery could be located near the center of the response field of the recorded neurons (as determined by microstimulation) or in the diametrically opposite position with a 2:1 ratio of occurrence. Only correct trials in which the target was presented in the recorded neurons' response field were included in analyses, and all analyses were performed separately for each neural population. Trials were further limited to only those with saccadic reaction times of greater than 100 ms to remove potential "cheat" trials. Unless otherwise specified, all analyses were performed using MATLAB 2019a (MathWorks) with custom code.

Spike times on each channel were first obtained offline using a voltage thresholding method. Each channel's spiking activity was then manually sorted into single units before continuing with analyses (using MKsort, a spike-sorting user interface, Ripple Neuro), and the low-dimensional representations of population activity are quite similar (see Figure 14). This is in line with previous work demonstrating that spike sorting has a negligible effect on the message of studies that focus on low-dimensional dynamics of population neural activity (Trautmann et al., 2019). Therefore, throughout the text we present results from spike-sorted neural activity but think of units (neurons) and multiunits (channels) as interchangeable. A total of 27 sessions were obtained and included in this study.

2.3.4 Dimensionality Reduction

In order to analyze the spiking patterns across the entire population, we utilized a dimensionality reduction method called Gaussian-process factor analysis, or GPFA (Yu et al., 2009). In short, this method converts spiking activity from a neural population into a lower-dimensional continuous "neural trajectory," where each dimension represents a weighted linear combination of neurons.

To perform GPFA, we used DataHigh (Cowley et al., 2013), a publicly available MATLAB code package for visualizing and reducing dimensionality of high-dimensional neural data. For a given session of laminar electrode data, all channels' spike times were first converted into spike trains aligned on target onset. Spike counts were grouped into non-overlapping bins of 20 ms width. Each observation includes data from one trial of the delayed saccade task beginning 200 ms before target onset and continuing through 200 ms post-saccade. This matrix is of size N channels x T time bins for each trial, with the latter dimension having a variable length. The GPFA algorithm returns a set of latent activity values summarizing the activity pattern across the population for each trial (matrix of L latent dimensions x T time bins for each trial). A cross-validation procedure was performed to determine the optimal number of reduced dimensions. The optimal dimensionality for each found via cross-validation was typically low (one to three). A final dimensionality of three was chosen for the sake of consistency across sessions and for ease of visualization in a 3D state space, as described next.

2.3.5 Defining Subspaces and Computing Proximity

The term "subspace" does not have a widely agreed-upon definition; some groups call each factor returned by dimensionality reduction a subspace in which latent neural activity could be varying (e.g., Kaufman et al., 2014) while others define new axes and focus on the variable activity along that dimension (e.g., Kobak et al., 2016; Libby & Buschman, 2021) or project activity onto an axis, plane, or hyperplane contained within a higher-dimensional space (e.g., Aoi et al., 2020;

Semedo et al., 2019). Here, we more loosely define a subspace as a distinct region in a low-dimensional state space occupied by neural activity during a specified condition (as in Churchland et al., 2006). The two main conditions here are "visual," or 100 to 200 ms after target onset (around the time of the putative visual burst) and "motor," or 120 ms to 20 ms before saccade onset (around the rising phase of the putative motor burst, and not including activity from after the latest time it is likely related to saccade initiation; see (Gandhi & Keller, 1999; Jagadisan & Gandhi, 2017; Miyashita & Hikosaka, 1996; Smalianchuk et al., 2018). A baseline condition was also defined, and this includes activity from 100 ms before target onset up to the time of target onset. These comprise the visual, motor and baseline subspaces, respectively.

For Figure 5, linear discriminant analysis (LDA) was performed to find the 2D projection that best separates visual and motor subspaces from each other. For each session, a two-class ("visual" and "motor" categories) linear discriminant classifier was trained and tested using a 10-fold cross-validation procedure. Only sessions for which the population activity was sufficiently separable between the visual and motor epochs (>70% classification accuracy rounded to the nearest integer, see Figure 5D) were further analyzed, leaving 22 sessions that met this criterion.

Given that the visual and motor activity (cyan and orange points, respectively) form distinct subspaces, we can use these as reference distributions against which to compare activity from time points throughout the course of each trial. We utilized a measure known as the proximity index, introduced by Dekleva et al., 2018. In short, the proximity index is a probabilistic measure that indicates the relative likelihood that a point of activity is closer in state space to a particular cluster than any other comparison cluster. For a single time bin of latent activity S, its proximity to the visual or motor cluster (VPI or MPI, respectively) is given by:

$$Proximity(S, \{C_i\}) = \frac{P(D_M(S, \{C_i\})|i)}{\sum_{j=1}^{3} P(D_M(S, \{C_i\})|j)}$$

where $\{C_i\}$ is the cluster of latent activity points during one of three reference conditions, denoted by i and j (here: visual, motor, or baseline) and $D_M(S, \{Ci\})$ is the Mahalanobis distance between the point S and cluster $\{C_i\}$. The VPI is formed when i =visual activity (and j =motor and baseline) and likewise, the MPI is formed when i =motor activity. The VPI and MPI are normalized to the range $[0, 1]$.

Since visual and motor proximity indices must be computed separately and result in two yoked values, we defined a visuomotor proximity index (VMPI) that can range from -1 to $+1$ and gives the relative proximity value of a point of activity to either the visual (-1) or motor $(+1)$ subspace:

$$VMPI = \frac{MPI - VPI}{MPI + VPI}$$

For Figure 6, the VMPI was computed for all non-overlapping 20 ms bins of latent activity throughout the time course of each trial. Activity from the baseline condition was treated as a third cluster to allow for the possibility of delay period activity existing in a completely different subspace than either the visual or motor subspaces, although proximities to this cluster are unimportant and hence not shown. It is also important to note that the absolute VMPI value ranges are inconsistent across populations (e.g., Figure 8A), but this does not matter for our study. Only the dynamics in the VMPI trace over the course of the delay period are of interest; thus, in Figure 8B-C and Figure 9B-C each population's VMPI trace was mean-subtracted to allow for a better comparison of sensorimotor transformation across populations.

2.3.6 Detecting Microsaccades and Aligning Proximity to Microsaccade Onset

All microsaccades that occurred during the delay period of each trial were detected offline. A 20 ms moving average of the eye velocity was taken, and a speed threshold of 5 to 15 deg/s was applied depending on noise level in the eye position signal. Saccades greater than 2 degrees in amplitude were rejected. Individual trials were manually evaluated to confirm correct automatic detection. Sessions were included in the following analysis if there were at least 20 trials in both the "one or more microsaccades" and "no microsaccades" conditions. One monkey (BB) did not consistently produce microsaccades during the delay period, as we have reported previously (Jagadisan & Gandhi, 2016). Hence, we could only include data from one session for this monkey using the above criteria. Monkeys BL and SU had five and eight sessions that met the above criteria, respectively, for a total of 14 sessions included in this set of analyses.

To determine the effect of microsaccades on the population activity pattern, we aligned the VMPI to microsaccade onset for trials in which at least one microsaccade was detected during the delay period. As a control analysis, we also aligned the VMPI to a pseudo microsaccade onset time for trials in which there was no microsaccade detected. For each trial, this alignment time was created by selecting a random time from the distribution of microsaccade onset times in trials with a microsaccade (Figure 9A-B). As in Figure 8, each population's trace was mean-subtracted to better compare trends across sessions.

2.3.7 Computing Relationship Between Population Activity Patterns and Behavioral Metrics

To examine whether the position in state space is related to the end behavior (i.e., saccade), we computed the correlation between the VMPI value at the animal's go cue and the eventual saccadic reaction time (RT) on that trial. We also asked if the position of the activity even leading up to the go cue was correlated with the end behavior. For this analysis, we worked backwards to compute the correlation coefficient of every 20 ms bin of activity with the saccadic reaction times on their respective trials. Values were tested for significance using a Wilcoxon rank sum test.

We also employed a similar approach developed by Afshar et al., 2011, that only utilizes information about the putative motor subspace rather than the visual subspace. In this framework, one can ask whether the distance traveled along the mean neural trajectory at the end of the delay period (equivalently, at the time of the animal's go cue), correlates with the saccadic RT on that trial. These methods have been described previously and were followed as closely as possible. In short, on a single trial, a vector of spike counts across the population starting at the time of go cue and going forward some short time in the future (dt=100 ms) is created. This vector is projected onto the vector created by mean values across all trials to obtain a projection value α (see Figure 10 inset). A correlation coefficient value between saccadic reaction time and this projection value was obtained for each session. To compute the correlation between activity prior to the go cue and the end behavior, we used the same value of dt (100 ms) but worked backwards to compute the median correlation coefficient of every 20 ms bin of activity with the RTs on their respective trials, as in the VMPI – RT correlation analysis. Of note, unlike all other results presented in this paper, this analysis was performed on spike-sorted but not dimensionality-reduced neural activity for a more direct comparison of findings across brain areas.

2.4 Results

In this study, we aimed to characterize delay period activity as exhibiting a visual- or motor-like signal and explored the relationship between this activity and the motor behavior to better understand the neural correlates of sensorimotor transformation in the SC. Neural activity from small populations of neurons was recorded simultaneously with multi-contact laminar probes traversing the dorsoventral axis of the SC as rhesus monkeys performed a delayed saccade task (Figure 3A-B). Data are examined from 27 recording sessions and limited to the subset of trials for which the visual stimulus and subsequent saccade were directed near the center of the response field. For electrode penetrations orthogonal to the SC surface, as we used here, all neurons recorded in a single session have comparable response fields, and care was taken to drive the electrode to the intermediate layers to capture mostly visuomotor neurons (i.e., those having both visual- and a motor-related increases in activity). A 24-channel laminar electrode was used for 15 of the sessions and a 16-channel electrode for the other 12 sessions. Spike sorting was performed and resulted in 12.3 (±3.3, range [7,19]) neurons per track, or population, on average. Across all 27 sessions, a total of 331 neurons were recorded.

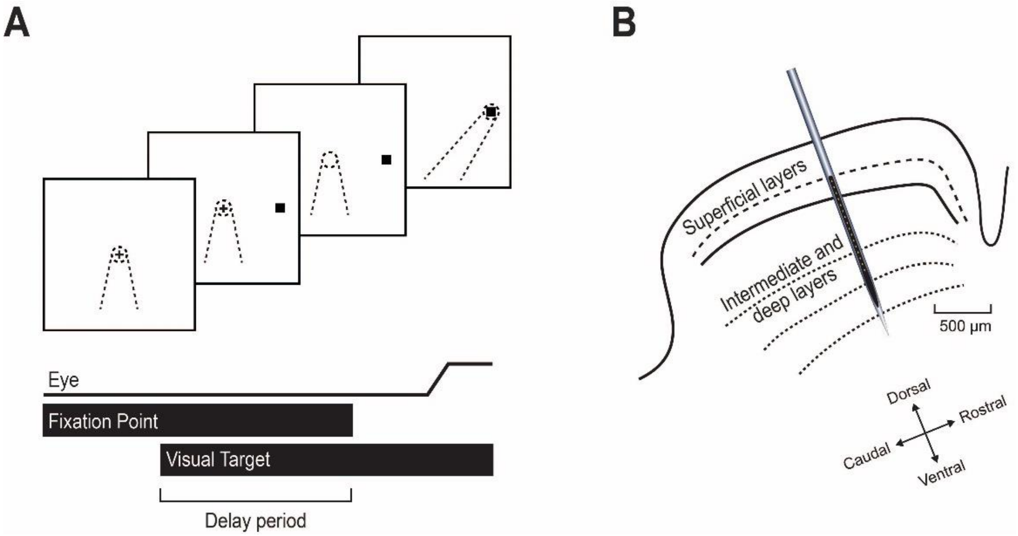

Figure 3. Schematic representation of behavioral task and neurophysiological preparation used in this study. Standard event sequence in the delayed saccade task. Top – The animal fixates on a central point (shown as a plus sign in this illustration). A target (black square) appears in the periphery for a variable "delay period." The dynamics of sensorimotor transformation were analyzed during much of the delay period. The disappearance of the fixation point acts as the go cue to generate a saccade toward the target. Center of gaze is depicted in all snapshots as a dotted cone. Bottom – A typical timeline of key events in a single trial of this task. B. In each experiment, a linear multielectrode array with 16 or 24 recording contacts was acutely inserted orthogonal to the SC surface along the dorsoventral axis to obtain a neural population representation as monkeys performed the delayed saccade task. Figure adapted from Jagadisan & Gandhi, 2022.

2.4.1 Visual and Motor Subspaces are Separable

We started by plotting for all contacts the average spike density profiles aligned on target and saccade onsets. We also separated the delay period from the transient visual burst, as shown for one session in Figure 4A. We then applied GPFA (Yu et al., 2009) to compute the latent activity patterns during these epochs (Figure 4B). For most datasets (23/27), including this one, the top 3 factors accounted for at least 95% percent of the variance in the spike density profiles (Figure 4D).

Thus, we limited our analysis to 3 dimensions, which also facilitated visualization. Moreover, instead of plotting the factors as a function of time, the low-dimensional activity can be illustrated in a three-dimensional state space, in which a single point denotes activity across the population taken from a 20 ms window from one trial (Figure 4C). This framework, on which we base our first set of analyses, allows an assessment of the regions, or “subspaces,” where the activity resides during the various epochs of the trial.

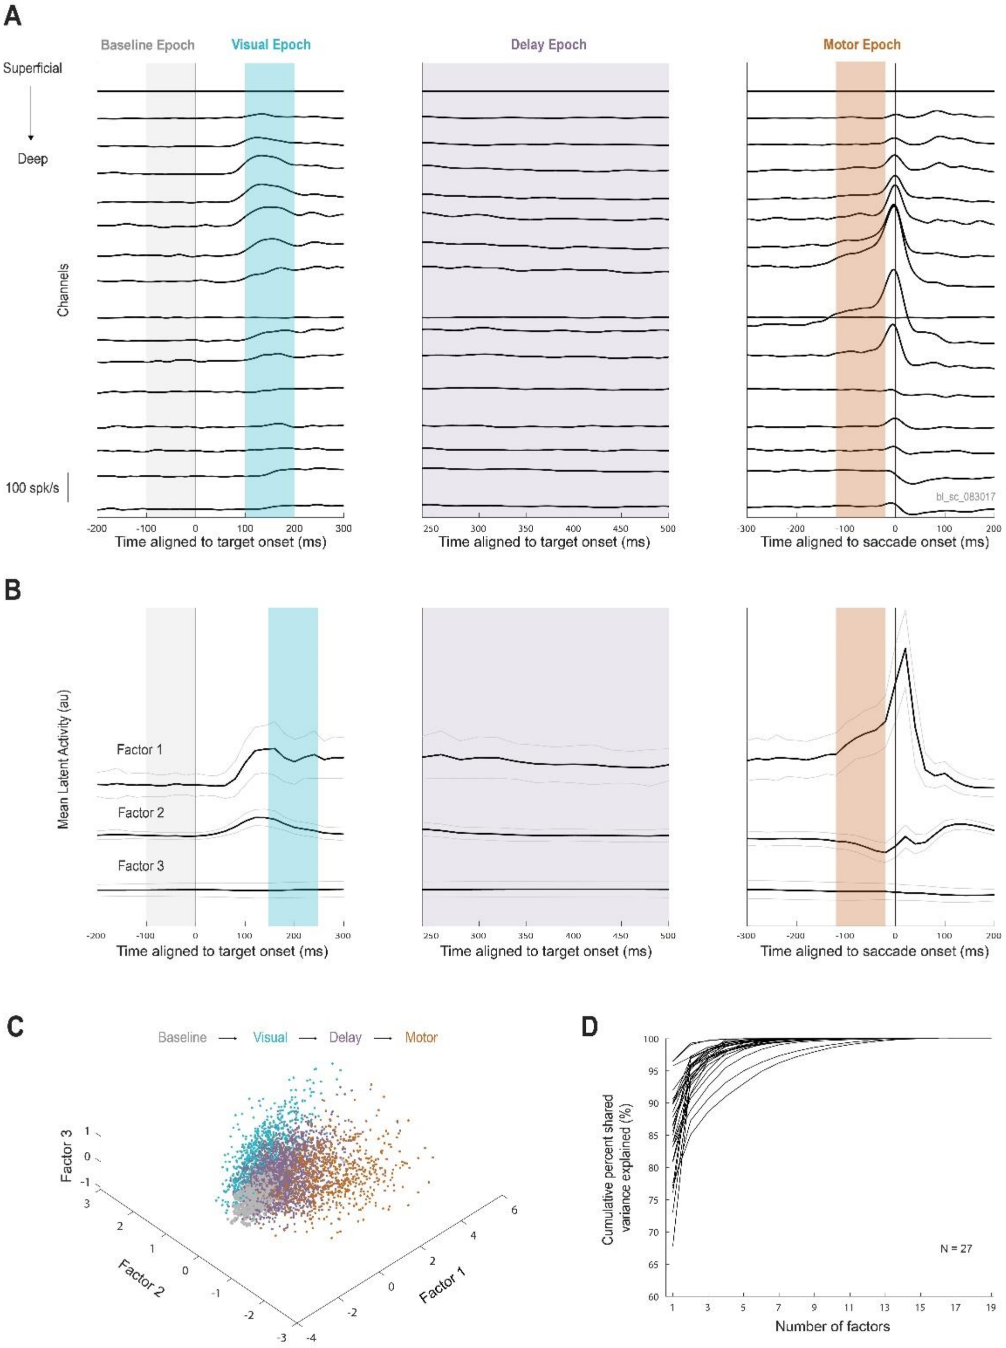

Figure 4. Analysis of neural population activity in a state space framework allows for an evaluation of subspace separability.

A. Trial-averaged firing rates across electrode depth are shown for one example session. Multiunit activity on each of 16 channels is plotted aligned to target onset (left and middle) or saccade onset (right). Each epoch window is defined by a vertical rectangle (baseline in gray, visual in cyan, delay in purple, and motor in

orange). Delay period activity was defined to start 240 ms after target onset, by which time the transient visual response had subsided. B. Latent population activity after dimensionality reduction using Gaussian Process Factor Analysis (GPFA) for the same example session after spike sorting into 12 single units. In each of the three panels, each trace is the trial-averaged (± one standard deviation) latent activity magnitude, plotted using the same alignment and epoch definitions as in (A). C. Latent activity represented in state space for the same example session. Latent activity during each of the four colored epochs are plotted as three-dimensional data points (each 20 ms bin has a magnitude along Factors 1, 2, and 3). Each dot represents the summary of the population activity pattern in a single 20 ms window; thus, a single trial contributes multiple points, even within the same epoch. D. Amount of covariability across neurons explained by lower-dimensional models compared to the full GPFA model. Each session is represented by a single trace. The majority of sessions have a high amount of shared variance explained by only one to three factors; thus, we retain the first three latent dimensions for each session.

In order to justify our comparison of delay period activity against two sets of activity, we first need to demonstrate that the activity patterns produced during the visual and motor epochs are distinct. Figure 5A-C shows the separability of the visual (100 to 200 ms after target onset) and motor (120 to 20 ms before saccade onset) latent activity patterns for three example datasets. By eye, the two subspaces are highly separable for these populations, and this separability was confirmed using linear discriminant analysis classification (Figure 5D). Across sessions, the mean classification accuracy was 82.7% (±11.3%), significantly above chance level of 50% (one-tailed t-test). This accuracy was not significantly different when considering multiunit activity (i.e., prior to spike sorting, see Figure 14). Our subsequent analyses required a high level of separability between the visual and motor subspaces. We used a minimum classifier accuracy of 70% as our cutoff criterion, which reduced our yield to 22 datasets.

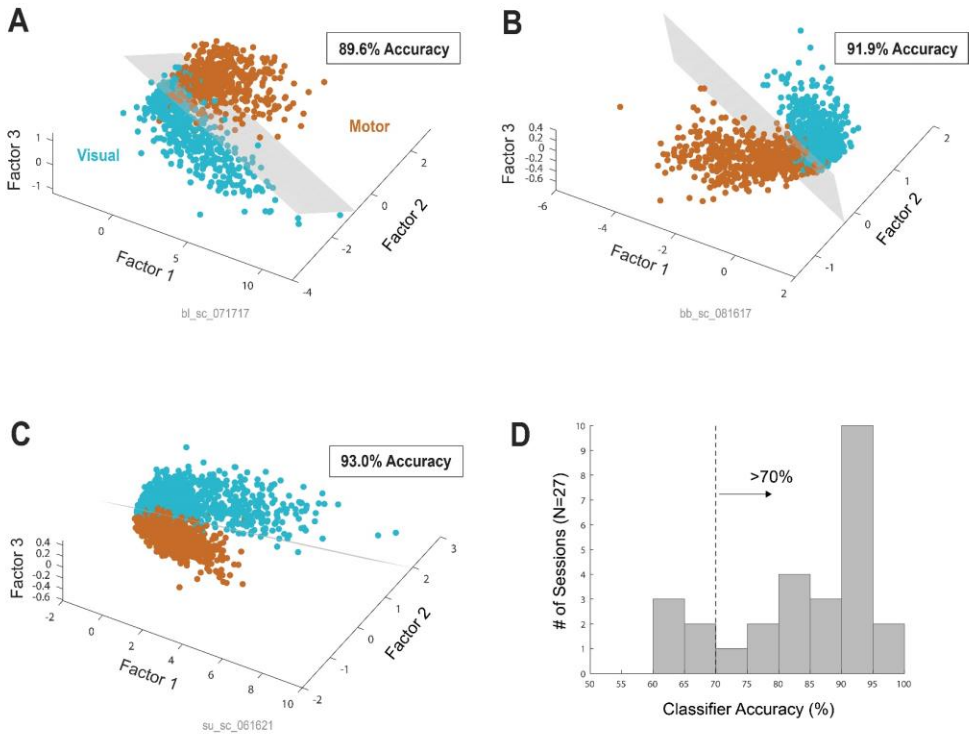

Figure 5. Visual and motor activity are linearly separable in state space for SC neural populations. A. Subspaces formed by latent activity patterns during the visual (cyan) and motor (orange) epochs for one example session. Time windows used for both epochs and a description of each data point are described in Figure 4. The gray shade indicates the plane of maximal separability as determined by linear discriminant analysis (LDA). The cross-validated classification accuracy of visual and motor points for this example session is also reported. B-C. Same as in (A) but for two additional example sessions. D. Histogram of linear discriminant classifier cross-validated accuracy in distinguishing visual from motor patterns across all 27 sessions. Only sessions with accuracy values of 70% or better were used in analyses that assume high separation between visual and motor subspaces, leaving 22 sessions for future analyses.

2.4.2 A Gradual Evolution from Visual to Motor Subspace Occurs During Sensorimotor Transformation

Once we established that visual and motor activity subspaces are separable, we wanted to examine the evolution of latent activity patterns throughout the delay period to determine if there is a consistent trend in the activity pattern from a visual-like representation to a motor-like representation. We utilized a Visuomotor Proximity Index measure (VMPI, see Methods) to compute the similarity of the population activity pattern during small windows of time to the visual and motor subspaces (Figure 6). As expected, the VMPI is very close to the visual subspace during the visual epoch (cyan shade), to the motor subspace during the movement epoch (orange shade), and in-between during the delay period (purple shade).

To evaluate the evolution of population activity in sensorimotor transformation, we plotted the VMPI during the delay period for all sessions before (Figure 8A) and after (Figure 8B) subtracting the mean of the first ten time bins for each session. These traces were then averaged to compute the session-averaged VMPI shown in Figure 8C. This trace reveals a slow, ramp-like progression of activity toward a motor-like representation as time in the delay period progresses. The monotonic trend was small but highly statistically significant ($p<0.001$, Mann-Kendall test), suggesting that the neural activity pattern slowly and systematically drifts toward a motor-like representation. The evolution of VMPI was highly variable across trials, exhibiting unique but noisy dynamics on individual trials (see two example sessions in Figure 7). Therefore, this monotonic trend only became evident when trial-averaging.

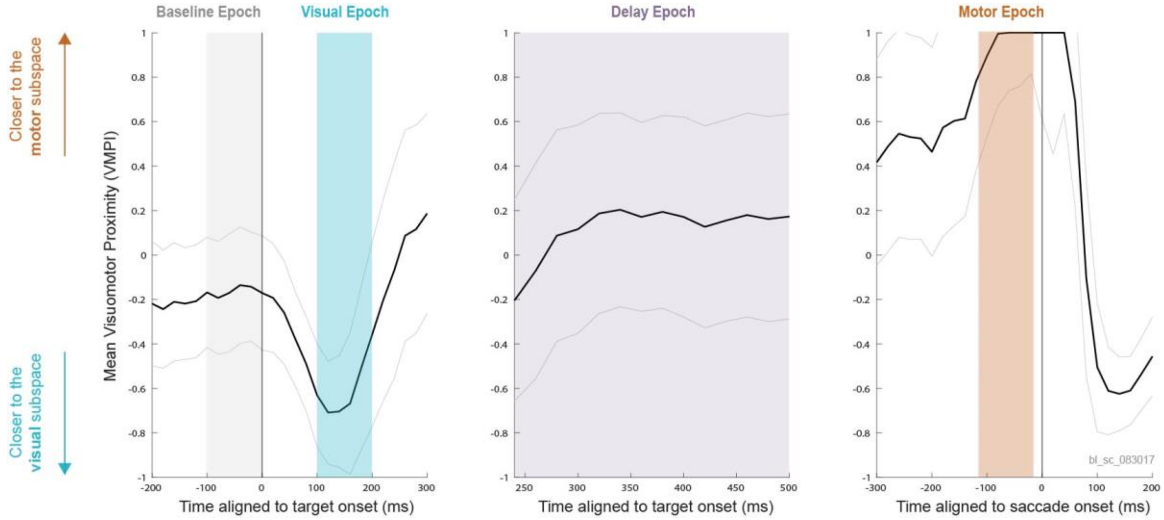

Figure 6. A Visuomotor Proximity Index (VMPI) can characterize the evolution of sensorimotor transformation.

Mean (± one standard deviation) VMPI values across trials for the same example session as Figure 4. The three panels also obey the same alignment and color scheme. A value closer to +1 indicates similarity of activity to a motor pattern and a value closer to -1 indicates that activity is more similar to the activity pattern produced during the visual epoch. VMPI values are by design limited to the range [-1,1]. The evolution of VMPI across the delay period (purple shaded rectangle) give insights into the representations of SC population activity between the visual and motor epochs.

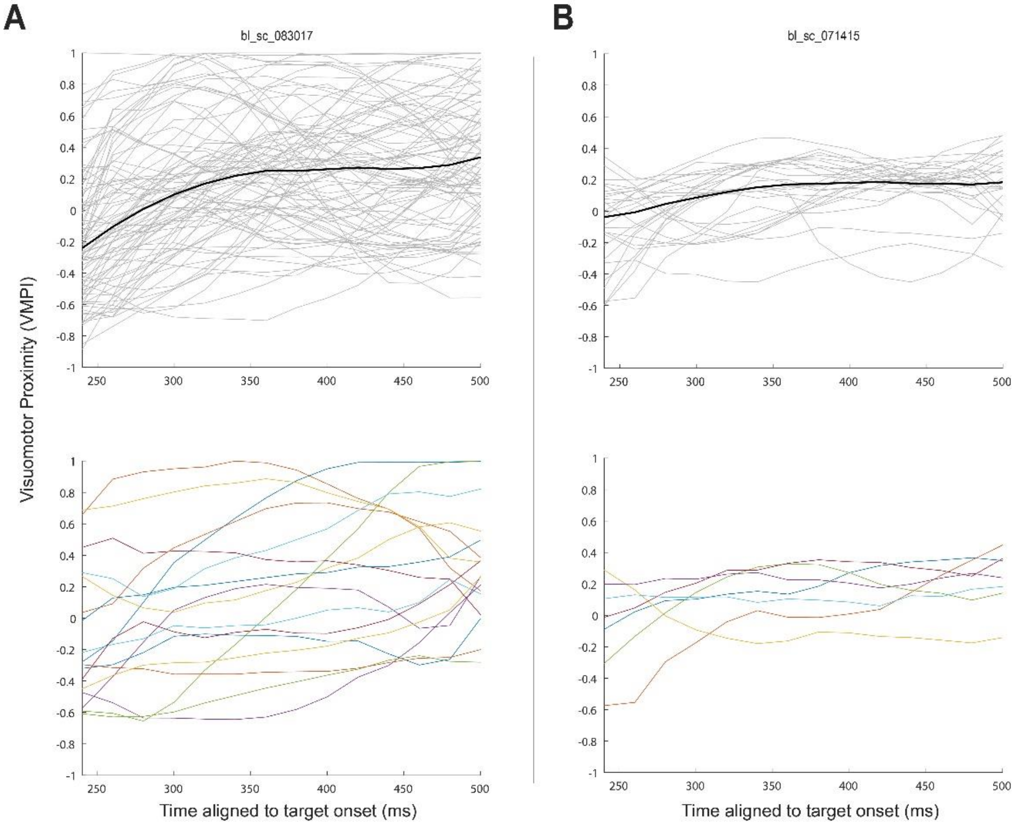

Figure 7. Single-trial delay period VMPI dynamics are highly variable.

A. TOP: VMPI values on individual trials (gray) of an example session (same as Figure 3 and Figure 6) after removing trials in which a microsaccade was made during the delay period. The across-trial median trace is shown in black. All traces have been smoothed with a 5-point moving average filter. BOTTOM: Individual trials from the same example session, subsampled from the entire pool of no-microsaccade trials and individually colored to highlight the across-trial variability in the VMPI trace. B. Same as (A) but for a second example session. Here, the range of VMPI values around the across-trial median is much smaller, yet a similarly broad range of dynamics is observable in single trials.

2.4.3 Microsaccades Transiently Revert Delay Period Activity Toward Visual Subspace

Microsaccades are rapid eye movements with kinematics resembling those of larger saccades but characterized by their small magnitude (less than ~2 degrees) and are frequently observed during fixational periods, including the delay period studied here (e.g., Hafed et al., 2015;

Jagadisan & Gandhi, 2016; Peel et al., 2016). Therefore, we questioned if microsaccades that occurred during the delay period produced perturbations in the neural activity pattern that impacted the monotonic trend observed in Figure 8C. Microsaccades were detected offline (see Methods), and for sessions with a sufficient number of trials with microsaccades (N=14), trials in which at least one microsaccade occurred at any point during the delay period were separated from those trials in which no microsaccade was made.

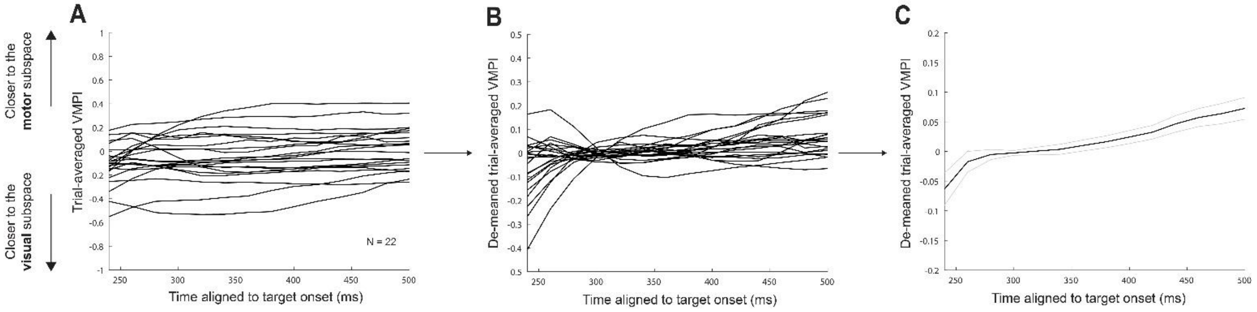

Figure 8. Trial-averaged neural activity slowly drifts from a visual- to a motor-like pattern. A. Trial-averaged Visuomotor Proximity Index (VMPI) value across the delay period for the 22 sessions with >70% accurate separability between the visual and motor clusters. Each session is represented as a single trace. B. Same as in (A) but with VMPI traces de-meaned for each session separately based on the first ten time bins. This allows us to average the VMPI traces across sessions to generate an across-session representation of the evolution of activity from the visual to motor subspace throughout the delay period, as shown in panel (C). There is a small but highly statistically significant, monotonic trend (p<0.001, Mann-Kendall test) from a more visual- to a more motor-like pattern throughout time in the delay period, indicative of a sensorimotor transformation signature.

Interestingly, when we aligned the VMPI value on microsaccade onset (microsaccade trials) or a semi-random delay period time (non-microsaccade trials, see Methods), we found a large and consistent effect, as illustrated in Figure 9A for an individual session's data and Figure

9B as an across session average. Beginning approximately 50 ms after microsaccade onset, the VMPI deviated toward the visual subspace, indicating a pronounced shift toward a visual-like representation that resolved over roughly the next 100 milliseconds. This aligns well with the idea that microsaccades serve to refresh information about the visual stimulus in the SC by jittering the target location on the retina (Khademi et al., 2020). However, at no time in the delay period was there a significant difference between the VMPI values for session-averaged microsaccade and no-microsaccade conditions (Figure 9C), indicating that the transient microsaccade-induced perturbation of the activity pattern on individual trials did not affect the slow, systematic drift observed on average throughout the delay period.

We decided to "zoom out" and see if this visual-like signature following a microsaccade was observable not only at the population level, but at the single neuron level. Figure 9D shows the trial-averaged firing rates of individual units aligned to microsaccade onset for two example sessions. For the session on the left, the more superficial units – those that typically exhibit a visual burst – clearly increase their firing rates following a microsaccade, consistent with the population-level analysis. This firing rate modulation of superficial units is nowhere near as pronounced for the example session on the right, yet the population-level measure we employed (i.e., VMPI) was able to easily pick up on a change in representation following a microsaccade (compare VMPI in Figure 9A with same session's individual neuron dynamics in Figure 9D right panel). We probed this effect further by correlating across all neurons and sessions the post-microsaccade firing rate with the visual burst evoked when a stimulus is presented in the response field (Figure 9E, R^2=0.64), the motor bursts generated for a saccade to that location (Figure 9F, $R^2 \approx 0$), and each neuron's traditional visuomotor index (see Methods; Figure 9G, R^2=0.34). The moderately strong relationship between visual activity and activity following a microsaccade confirms that the transient reset toward the visual subspace seen in Figure 9B in large part arises due to a strong transient increase in the activity of neurons in the population that have a visual burst.

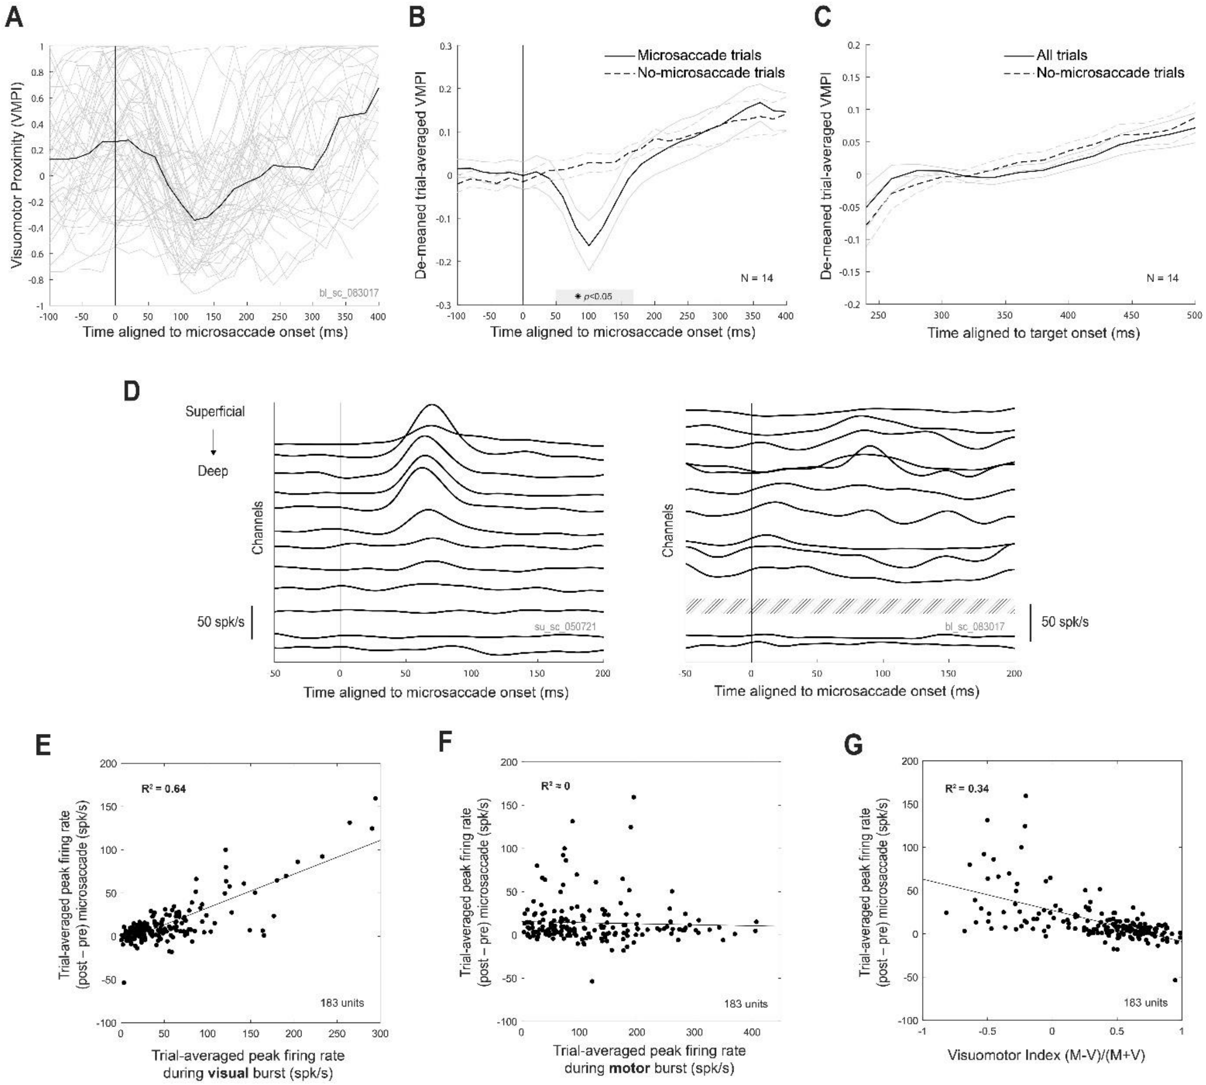

Figure 9. Single-trial population neural activity transiently reverts towards a visual pattern after a microsaccade.

A. VMPI values on individual trials (gray) of an example session (same as Figure 4 and Figure 6) in which one or more microsaccades were detected during the delay period, with the median trace shown in black. VMPI values are aligned to microsaccade onset time, regardless of the absolute time in the delay period the microsaccade occurred. On average, the VMPI value dips toward a visual-like pattern ~50 ms following a

microsaccade. B. Session-averaged, de-meaned VMPI values (mean in black ± one standard error of the mean in gray, as in Figure 8C) aligned to microsaccade onset (solid lines) or pseudo-microsaccade onset on non-microsaccade trials (dashed lines, see Methods). The population activity pattern significantly and transiently deviates from the pattern produced on non-microsaccade trials starting around 50 ms and returns to match the pattern produced on non-microsaccade trials by 170 ms (Wilcoxon rank sum test, $p<0.05$). C. Same setup as in Figure 8C, with de-meaned, session-averaged VMPI values throughout the delay period plotted for two groups of trials – all trials (solid lines) and trials in which no microsaccade was detected during the delay period (dashed lines). The two traces never significantly differ from each other (Wilcoxon rank sum test), demonstrating that the presence of microsaccades on some trials does not affect the interpretation of Figure 8C. D. Trial-averaged firing rates of individual neurons aligned to microsaccade onset for two example sessions. Each unit is offset vertically based on the channel on which it was recorded (hence, the example session on the right has two units shown on one channel and no units shown on another, as indicated by diagonal shading). For many sessions, the superficial neurons burst following a microsaccade, although in some sessions – such as the example session on the right – it is difficult to appreciate any activity change following a microsaccade when looking at individual neuron firing rates. Still, even for these sessions, there is a population-level transient shift in representation toward a visual-like pattern (see same example session in (A)). E. Neurons with a visual response exhibit increased activity levels following a microsaccade. Each point represents a single neuron from a single session. F-G. Same setup as in (E) but for the relationship between peak activity around saccade onset (F) or visuomotor index (G) and the peak firing rate following a microsaccade. There is no correlation between the level of motor-related and microsaccade-related firing rates, and there is a lower correlation between the relative visual and motor properties of neurons (using Visuomotor Index, VMI, as a proxy) and microsaccade-related activity than when only the strength of the visual burst is considered (i.e., E).

2.4.4 Sensorimotor Transformation Process is Predictive of Reaction Time

We show above that on average, SC activity slowly drifts toward a motor-like representation throughout the delay period. This prompted us to ask if the magnitude of this drift on each individual trial was related to the animal's ability to rapidly initiate an eye movement on that trial. We hypothesized that the higher the VMPI at a given time in the delay period (and therefore the larger the drift), the less time it would take the monkey to initiate the saccade after a cue to initiate movement ("go cue") was given on that trial. Figure 10A shows the correlation coefficients between the VMPI value at time windows leading up to the animal's go cue and the eventual saccadic reaction time (RT). At the time of the go cue to make an eye movement (time=0 ms, rightmost point of Figure 10A), the VMPI value was significantly correlated with the eventual RT, supporting the idea that the amount of drift on an individual trial by the end of the delay period is predictive of the animal's ability to initiate a movement. This relationship held for all time points during the delay period leading up to the go cue (time range of -360 to 0 ms, Figure 10A). No other saccade metrics (i.e., amplitude, peak velocity, endpoint error) were found to be correlated with RT (Figure 11A-B), nor might we expect them to be (Hafed, 2021).

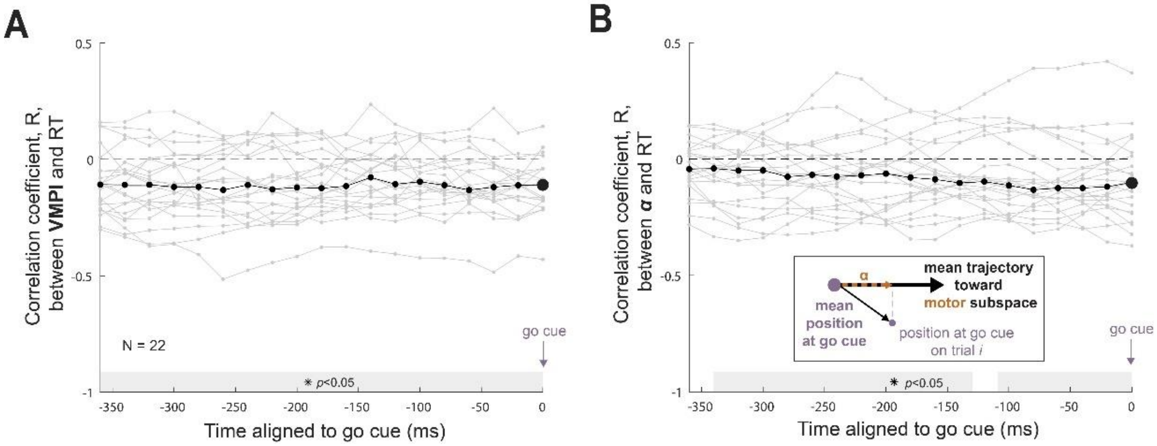

Figure 10. The single-trial state-space position of activity is correlated with that trial's saccadic RT even long before the go cue.

A. Across-trial correlation coefficient between the VMPI value in a single time bin relative to that trial's go cue time and the eventual saccadic reaction time (RT) on that trial. Each gray trace represents the across-trial correlation coefficients for a single session, with the across-session median trace shown in black. Time bins in which the median correlation coefficients were significantly below zero ($p<0.05$, one-tailed Wilcoxon signed rank test) are shaded along the x axis in gray. Even long before the go cue, the state space position of activity (as computed via VMPI) is correlated to a behavioral metric. B. Same as in (A) but for correlations between a single trial's projection value α and that trial's saccadic RT. Each projection value α represents the distance traveled along the trial-averaged neural trajectory toward the motor subspace by a certain time relative to the go cue on a given single trial *i*, with the method for finding α shown in the inset. Methodology was previously used in Afshar et al., 2011, and applied here to SC neural populations. This different method of computing state-space position of activity reveals a similar correlation to saccadic RTs, up to 340 ms before the go cue on average.

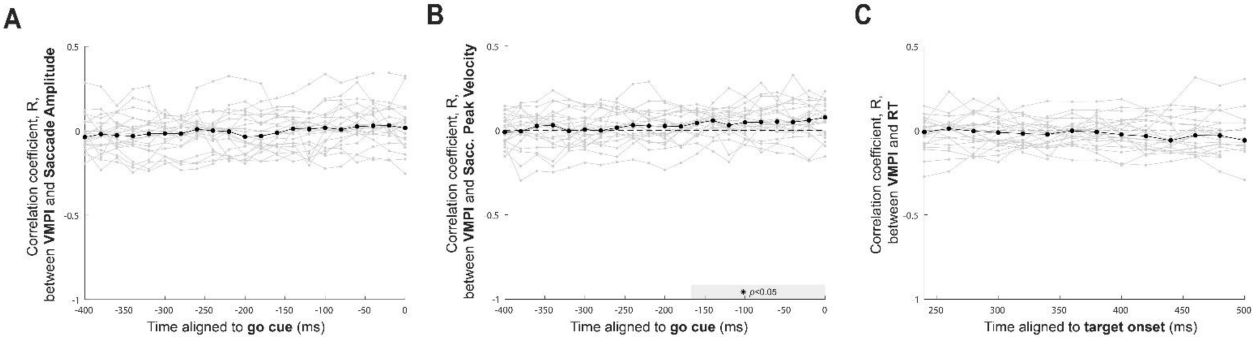

Figure 11. The relationship between VMPI and saccade metrics is variable.

A. Across-session median correlation coefficient, R, between VMPI and saccade amplitude, for times leading up to the go cue. Same conventions as in Figure 10. There is never a significant correlation between the two variables (Wilcoxon signed rank test). B. Same as in (A) but for correlations between VMPI and single-trial peak velocity of the saccade. Time bins in which the median correlation coefficients were significantly different from zero (p<0.05, Wilcoxon signed rank test) are shaded along the x axis in gray. A positive relationship between the two variables emerges approximately 160 ms before the go cue. C. Same as in Figure 10A but for correlations between the VMPI value aligned to the beginning of the delay period and single-trial RT. No significant correlations are observed (one-tailed Wilcoxon signed rank test).

When aligning single-trial VMPI values to the beginning of the delay period (Figure 11C), the position of the activity was not correlated to the eventual behavioral output, suggesting that the population representation likely drifts at relatively similar rates across trials. If so, then the drift would be bigger for the longer delay trials, which would be associated with faster reaction times. Indeed, the variable length of the delay period, in the context of our experimental paradigm, seems to account for the variable latency in the behavior (Figure 12A-B).

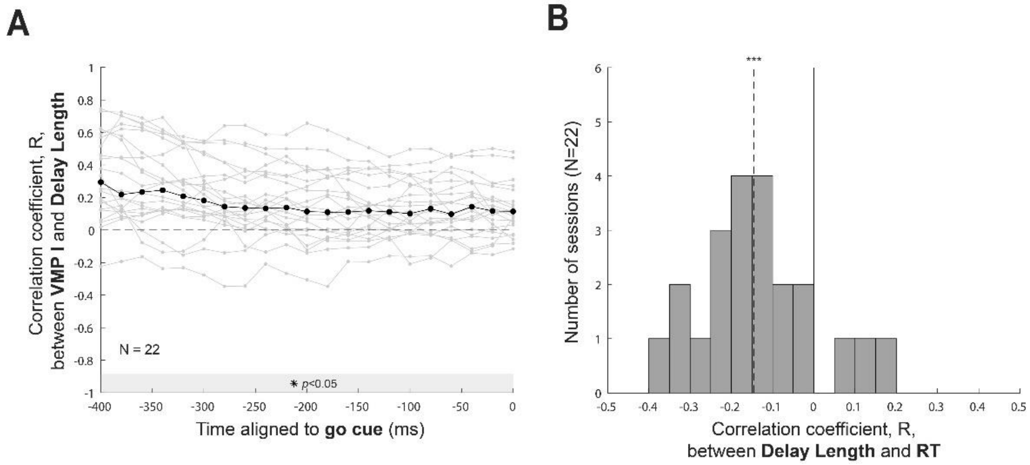

Figure 12. Delay lengths are correlated with both VMPI and RT.

A. The VMPI value is significantly correlated with the delay period length even 400ms before each trial's go cue time (p<0.05, Wilcoxon signed rank test). Same conventions as in Figure 10, with the across-session median correlation coefficients shown in black and individual sessions' correlation values shown in gray. B. Histogram of correlation values between each trial's delay period length and saccadic reaction time for the 22 sessions included in analysis of sensorimotor transformation. The across-session median correlation coefficient (-0.145) was significantly less than zero (p<0.0001, one-tailed Wilcoxon signed rank test), indicating an inverse relationship between delay period length and reaction time.

This finding conforms well to an existing theory of arm movement generation – the initial condition hypothesis (Afshar et al., 2011; Churchland et al., 2006) – in which the population activity pattern at the animal's go cue informs the latency of the reach initiation on that trial. Critically, however, the brain areas relevant to reach initiation (e.g., primary motor cortex and dorsal premotor cortex) do not have significant responses to visual stimuli. Because of the strong visual bursts exhibited by SC neurons, we wanted to ensure that the relationship between population activity pattern and latency of saccade initiation remains even when disregarding the visual-likeness of the pattern during the delay period. Therefore, we decided to employ an additional methodology (mirroring that of Afshar et al., 2011) that only considers the position and trajectory of activity patterns relative to the patterns underlying *motor* output rather than relative to activity during both the visual *and* motor epochs (as is the case with the VMPI metric). In short, two vectors are created – one that extends from the mean activity position at the time of the go cue to some short time later (100 ms) and another that extends to the actual position at the time of the go cue (or before) on an individual trial (shown in Figure 10B inset). The projection value α obtained by projecting the latter vector onto the former gives you a magnitude that can be thought of as distance traveled toward the motor subspace by a certain time in the delay period.

When applying these methods to SC population activity, we found that the correlation values closely matched both those applied previously to premotor cortex activity and those obtained through our VMPI – RT correlation analysis. Across many populations, the median correlation between projection value at the go cue and RT across trials was small but significant (Figure 10B, data points at t=0 ms), similar to the results of the VMPI – RT correlation analysis (i.e., Figure 10A). Also comparable was the significant correlation between projection value α and eventual RT even long before the go cue (Figure 10B, data points leading up to t=0 ms). Therefore, this relationship between the population-level activity pattern and movement initiation latency holds even when only considering the motor-likeness of the pattern.

2.5 Discussion

In this study, we sought to understand whether population activity in the SC systematically transitions from a visual-like pattern to a motor-like pattern throughout the delay period and found that on average, the activity pattern (as measured by the VMPI value) did indeed slowly drift from a visual-like to a motor-like representation. Notably, following a microsaccade, the activity pattern was characterized by a transient reset to a visual-like representation (Figure 13A). In addition, the amount of drift exhibited by the population at times leading up to the animal's go cue on individual trials was predictive of the latency at which a saccade could be initiated on that trial. On the other hand, neither the starting representation nor the amount of drift aligned to the beginning of the delay period were related to the eventual saccade latency. Together, these findings lead us to conclude that activity drifts at a relatively consistent rate across trials and sets the animal up for a shorter or longer latency saccade based on the amount of time the activity has had to drift on a given trial (Figure 13B). This study therefore provides new insights into the neural dynamics expressed within the SC during the delay period of a widely used behavioral task.

Our population-level results disclosed a gradual evolution from a sensory-like to motor-like representation that is consistent with the interpretations provided from previous studies on sensorimotor transformation in the SC (e.g., Lee and Groh, 2012; Sajad et al., 2020; Sadeh et al. 2020). If the SC is indeed involved in sensorimotor transformation, then we could have expected to observe a few possible alternate dynamics in the visual-like or motor-like representations throughout the delay period. First, activity patterns could oscillate between the visual and motor subspaces, balancing two needs – retaining information about the sensory stimulus and preparing for a movement. This is akin to the idea of maintaining simultaneous representation of multiple auditory stimuli (Caruso et al., 2018a). Second, activity patterns could lack a consistent trend toward either subspace across the delay period, potentially encoding features independent of sensory information or movement preparation (Kaufman et al., 2015). Third, activity patterns could exhibit a discrete switch, or step, from one representation to the other (Latimer et al., 2015). Instead, the slow evolution we saw in SC population activity patterns toward a motor-like pattern is evidence of a smooth and slow drift in representation between the times of sensation and action.

Furthermore, recognize that the subspaces labeled as sensory and motor are merely a reduced dimensional representation of the population activity. Importantly, they are agnostic to the reference frame represented by the neural activity. Finally, we note that the gradual transition was only observed when averaging across trials. Despite the many benefits of dimensionality reduction, neural activity on single trials is still inherently noisy, and thus the lack of quantifiable trend on individual trials does not decrease our confidence that the across-trial sensorimotor transformation trend is meaningful.

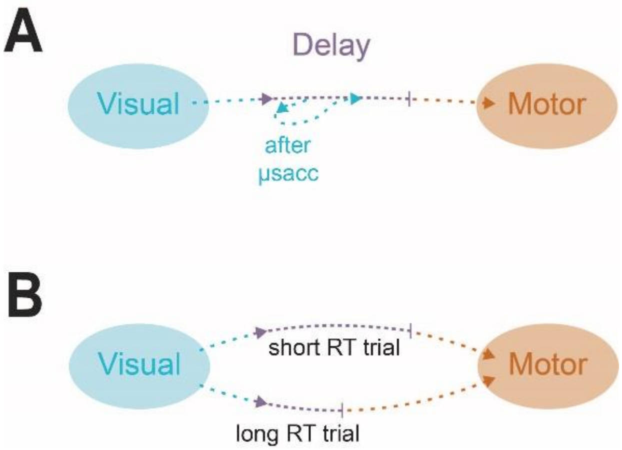

Figure 13. Schematics representing the evolution of SC population activity throughout the course of a standard sensorimotor task.

A. Characteristic time course of population activity during the delay period. On average, population activity slowly becomes less visual-like (visual subspace shown as a cyan ellipse) and more motor-like (orange ellipse) throughout time in the delay period (purple dotted line). Following a microsaccade, when it occurs, population activity transiently and strongly deviates its trajectory toward the visual subspace (cyan dotted line) before returning to the characteristic route toward the motor subspace. Vertical purple line denotes VMPI at go cue. B. Depiction of the population activity trajectories on microsaccade-absent trials with short (top) and long (bottom) saccadic reaction times. Activity during the delay period likely evolves at a comparable rate across trials, at least under the experimental conditions we used. However, the activity can continue its slow drift toward the motor subspace for a longer time on trials with longer delay period lengths (top) than on those with shorter delay period lengths (bottom), thereby only necessitating a limited distance to travel to the motor subspace once the cue to initiate the behavior is given (vertical purple line).

In one study, subpopulations of cortical oculomotor neurons categorized based on their relative firing rates during the visual and motor epochs were shown to exhibit unique temporal dynamics of this reference frame transformation, in contrast to the smooth and gradual transition observed when treating all neurons as a single population (Sajad et al., 2016). The VMPI measure we use implicitly takes into account the activities of all neural subtypes (i.e., visual, visuomotor, and motor) and produces one concise value, although the neurons recorded in this study were by and large visuomotor neurons. Research that specifically teases apart the individual contributions of neural subtypes to the time course of sensorimotor transformation may be required for a more complete understanding of the SC correlate of this behavioral phenomenon. Also not considered within this study are cognitive factors such as reward anticipation, arousal, and attention – factors which may in fact be multiplexed with encoded visual and motor signals present in the SC. Thus, the exact relationship of delay period SC activity patterns to phenomena other than visual processing and movement initiation is ripe for future investigations.

2.5.1 Does SC Delay Activity Resemble a Preparatory Signal?

The results shown in Figure 10A and B suggest that a relationship between the activity pattern and the reaction time (RT) of the eventual saccade is present long before permission is granted to initiate the movement. In other words, if the pattern of SC population activity drifts close to the motor subspace (i.e., has a high VMPI value) during the delay period of a given trial, the activity will take little time to evolve into a fully motor-like pattern after the go cue, resulting in a low-latency saccade. In the context of our task design, SC delay period activity seems to drift in representation at a relatively equal rate across trials (Figure 11C). It is on trials with longer delays that the population activity has extra time to evolve, and therefore continues to drift toward a motor-like representation (equivalently, an increasing VMPI) proportional to the delay period length, resulting in proportionally fast reaction times (Figure 12A-B).

The observation that the rate of drift is unmodulated from one trial to another could be reflective of the animal’s internal model of the expected delay period length distribution. To extrapolate, if the delay period length was known by the animal to instead be constant, the activity pattern may still drift at slightly different rates from trial to trial, which we posit would constitute preparatory activity and serve as a mechanism for movement initiation. Regardless, the schematic shown in Figure 13 is inclusive of both schemas.

The consistent rate of sensorimotor evolution from trial to trial leads us to consider whether this drift might act as a self-timing mechanism, indicative of perceived time elapsed in the delay period. Suppose the monkey employs a strategy in which he begins to expect a delay period length somewhere in between the shortest and longest delays previously experienced (as observed in human studies; Jazayeri & Shadlen, 2010). If he consistently plans to initiate an eye movement after this expected delay period length, he may receive the benefit of more frequently successful trials (and consequently, more frequent and more rapid rewards) since he has optimized the timing of his saccade to match his expectation. Neurons in the macaque thalamus (Tanaka, 2007) and lateral intraparietal cortex (Leon & Shadlen, 2003) have been shown to encode perceived time intervals. Perhaps the drifting representation we observe in SC populations is another signature of task timing. Further experiments might explore this concept and its validity.

Although we consider SC activity during the delay period to be preparatory in the sense that it is related to the enhancement or hindrance of rapid saccade initiation following the go cue, it does not have "motor potential." In the smooth pursuit system of the FEF, neural activity *was* found to have motor potential, with partially overlapping subpopulations contributing to both the preparation and execution of movement (Darlington & Lisberger, 2020). However, this does not seem to be the case for the SC, at least in the context of saccades. Previously, we have demonstrated that inhibition of the omnipause neurons during the delay period, which allows SC activity to travel to saccade-generating brainstem structures, is not sufficient to evoke a saccade (Gandhi & Bonadonna, 2005; Jagadisan & Gandhi, 2017). In addition, the lack of a burst and only a baseline-level “temporal stability” – that is, the consistency in the population activity pattern from one time point to another – enhance the argument against delay activity having motor potential (Jagadisan & Gandhi, 2022). We have also observed that it is typically the neurons with strong visual bursts rather than strong motor bursts that have sustained activity during the delay period (Massot et al., 2019). Therefore, it stands to reason that preparatory signals are likely encoded in the SC in dimensions orthogonal to those during movement (e.g., orthogonal potent-null subspaces; Kaufman et al., 2014).

2.5.2 The Relationship Between Microsaccades and SC Population Activity Patterns

Microsaccades produced during fixation serve to refresh the visual stimulus on the retina in order to combat a fading perception over time (Martinez-Conde et al., 2004). The neural circuit in SC suppresses vision during microsaccades (Hafed and Krauzlis, 2010), as it does during large amplitude saccades (Robinson and Wurtz, 1976). Following the movement, the nervous system responds to its visual environment by evoking activity in visually responsive neurons, although extra-retinal sources likely contribute as well. Indeed, we observed that microsaccades produced during the delay period consistently perturbed the sensation to action transition by transiently deviating SC population activity toward the visual subspace. The effect was strongest in the subset of neurons with a robust visual response. This modulation began roughly 50 ms after microsaccade onset and peaked another 50 ms later before rapidly meeting back up with the population activity patterns observed in non-microsaccade trials. Thus, the resurgence of visual activity likely reflects visual reafference following the movement (also see Khademi et al., 2020).

It is valuable to consider the various ways in which microsaccades generated during the delay period could have impacted the oculomotor system. For instance, the movement-related activity associated with microsaccade generation could have accelerated the overall SC output toward the motor subspace, resulting in reduced saccade latency – perhaps even triggering it prematurely before the go cue – and altered endpoint accuracy (Buonocore et al., 2021). Instead, we observed a rapid rebound and return of the system (VMPI trace, Figure 9B) to its original trajectory following a microsaccade. We interpret this reset to suggest that the gradual transition from a sensory to motor representation may be a network feature that is resistant to the effects of transient disruptions. As a whole, these observations lead us to conclude that microsaccades are a potential mechanism for engaging the network to produce a visual-like signal very similar to that elicited in response to the initial target appearance, but one that is compensated for quickly and robustly.

2.5.3 Low-Dimensional Geometry of SC Population Activity and its Skeletomotor Counterparts

One of our objectives was to extend to the oculomotor system the dynamical systems perspective of motor control that has been studied extensively in the skeletomotor system (Shenoy et al., 2013; Gallego et al., 2017). Studies of arm reaching that use this framework have given rise to multiple hypotheses for mechanisms of movement initiation. One such schema is the "optimal subspace hypothesis" (Churchland et al., 2006), which propounds that there is an optimal set of population activity patterns that allow for the generation of a goal-directed movement. The initial condition hypothesis (Afshar et al., 2011) builds on this framework by postulating that trials in which patterns that have traveled closer to the motor subspace by the time of the animal's go cue will have a faster reaction time (RT) than those in which the underlying neural activity has not traveled as far along the mean neural trajectory.

It might make sense for the oculomotor system to operate in a somewhat different manner than the skeletomotor system given the additional element – visual information – encoded within the SC and other oculomotor areas. However, even when considering this additional set of patterns exhibited by SC populations, we found that on trials in which population activity more closely resembles a motor-like pattern, the saccadic RT is significantly shorter (Figure 10A).

To establish a more direct comparison between SC activity and activity in its skeletomotor analogs, we also applied the methods of Afshar et al., 2011, to SC population activity during the delayed saccade task and found a comparable, significant correlation between the position of delay period activity and the saccadic RT (Figure 10B). Although the exact methodologies applied in Figure 10A and B are distinct, they address similar questions – primarily, is the similarity of neural activity during the delay period to motor activity related to the speed at which the movement can be initiated (i.e., eye movement or reaching movement RT). The findings reported in this study support the idea that the initial condition hypothesis is also valid for the oculomotor system.

The optimal number of dimensions needed to explain the across-trial shared variance of our acutely recorded neural populations was much lower than that described in studies using neural activity recorded in primary motor cortex (M1) or dorsal premotor cortex (PMd), for example (Churchland et al., 2010; Churchland & Shenoy, 2007). We conjecture that this is at least partially due to the homogeneity of each recorded population. As SC neurons are traditionally recorded along a dorsoventral axis, the topography of the SC yields populations in which each neuron has a similar response field, chiefly varying across electrode depth in the strength of their visual and motor bursts (Massot et al., 2019). Cortical areas like M1 and PMd yield much more heterogenous populations with respect to the spatial locations preferentially encoded by each neuron, and the dynamics underlying behavior are typically studied after grouping trials with multiple reach directions. However, since our recorded SC populations vary not in their preferred spatial location but rather in their visual and motor signal strengths, we limited our analyses to a single saccade direction so that in reducing the dimensionality of the data, the variability between visual and motor patterns would be brought to the forefront.

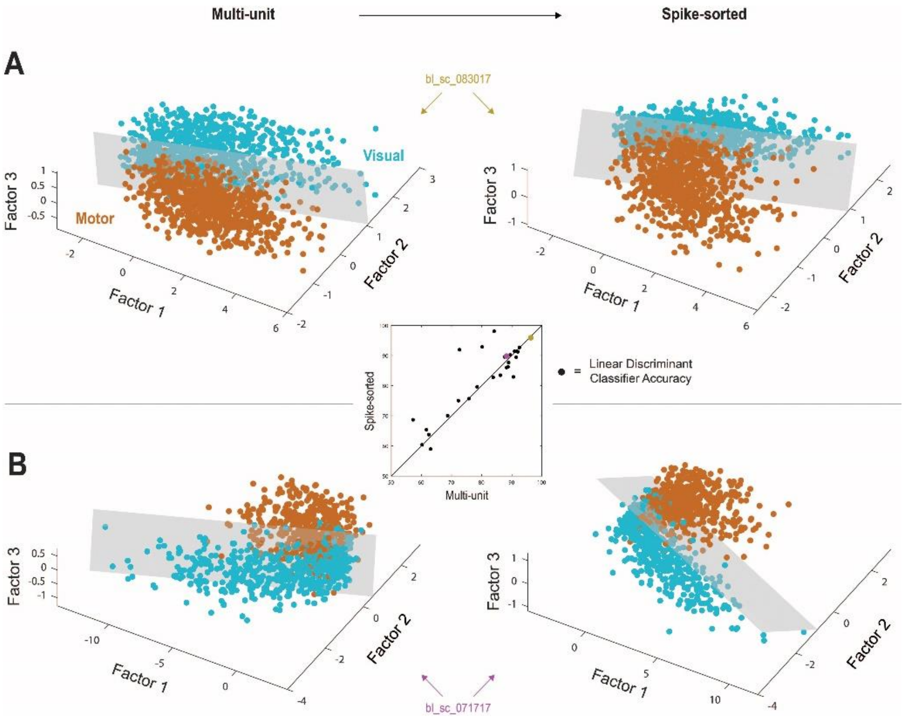

Figure 14. Spike-sorted and multiunit populations exhibit nearly-identical activity patterns throughout the delayed saccade task.

A. Low-dimensional representations of population activity patterns during the visual (cyan) and motor (orange) epochs for one example session before (left) and after (right) spike sorting. Although each dimension is not directly comparable across multiunit and spike-sorted populations, the subspaces formed by population activity in both cases are nearly identical. B. Same as in (A) but for a second example session. The subspaces formed by visual and motor activity before and after spike sorting have similar levels of separability. The exact position of each point of activity is unimportant to the comparison across epochs. Inset. A comparison of the visual and motor subspace separability obtained through LDA classification pre (x axis) and post (y axis) spike sorting for all 27 sessions. Accuracies for example sessions shown in (A) and (B) are colored in gold and magenta, respectively.

3.0 Contextual Effects on Population Representation of Sensory and Motor Responses in the Superior Colliculus

This work to be submitted as a manuscript with contributing authors: Ayar EC, **Heusser MR**, Bourrelly C, Gandhi, NJ.

3.1 Overview

Sensorimotor transformation is the process of first sensing an object in the environment and then producing a movement in response to that stimulus. For visually-guided saccades, neurons in the superior colliculus (SC) emit a burst of spikes to register the appearance of stimulus and many of the same neurons discharge another burst to initiate the eye movement. We investigated whether the neural signatures of sensation and action in SC depend on context. Spiking activity along the dorsoventral axis was recorded with a laminar probe as Rhesus monkeys generated saccades to the same stimulus location in tasks that require either executive control to delay saccade onset until permission is granted or the production of an immediate response to a target whose onset is predictable. Using dimensionality reduction and discriminability methods, we first verified that the subspaces occupied during the visual and motor epoch were distinct within each task. Surprisingly, the low-dimensional spaces spanned by sensation and action epochs were also differentiable across tasks. These results imply that cognitive processes associated with task requirements are multiplexed in SC population activity during both sensation and action, and that downstream structures could use this activity to extract cognitive context. They also suggest that the entire manifolds associated with sensory and motor responses, respectively, may be larger than the subspaces explored within a certain set of experiments.

3.2 Significance Statement

Sensorimotor transformation is a process that humans perform over 100,000 times a day, for example, when we look at or reach for objects of interest. Many areas of the brain register a sensory stimulus and convert the stimulus-related information into an appropriate motor output. Deficits in sensory and/or motor processes are implicated in neurological dysfunctions like attention deficit disorders. We investigated how the context of behavioral tasks impacts patterns of population activity during sensation and action. Dimensionality reduction techniques reveal that context can impact both the sensory- and movement-related transient activity patterns in the population of superior colliculus neurons. This insight is not readily identifiable in single unit recordings.

3.3 Introduction

Envision a stopped car in a turn-only lane at an intersection. When the traffic signal turns green, the driver will either make the turn immediately or wait until oncoming traffic and pedestrians in the crosswalk clear. In both cases, the primary sensory input and motor output – traffic light turning green and making the turn, respectively – are categorically identical, but cognition contributes to the overall process. Executive control clearly impacts when the movement is produced, but it is not apparent whether the neural signals underlying sensation and action are also altered. Such conditions can be simulated in the laboratory environment and combined with neurophysiological experiments. Accordingly, we relied on the wealth of knowledge in the oculomotor system to study contextual impact on the transient sensory and motor representations in the superior colliculus (SC) during visually guided saccades.

Cognitive context can be manipulated through two behavioral tasks that are akin to the aforementioned scenarios. The visually-guided delayed saccade task (which we will refer to as the “delay task” throughout the text) requires executive control to withhold the eye movement until another cue grants permission. This condition prolongs the temporal delay between the sensation and action aspects of the task. The “gap task,” in contrast, requires an immediate action. It creates a predictable situation in which movement preparation can precede stimulus onset, which can reduce further the reaction time (Saslow, 1967; Dorris and Munoz, 1995; Sparks et al., 2000).

Situated at the roof of the brainstem, the SC is a topographically organized, laminar structure that encodes sensory and motoric information along its dorsoventral axis (Wurtz et al., 2001; Gandhi and Katnani, 2011; Basso and May, 2017; Massot et al., 2019). A population of neurons will emit a burst of spikes when a visual stimulus is presented in their response fields. As this initial response is largely exogenously driven, it reflects “sensation” or registration of the stimulus. Just before saccade onset, an ensemble of SC neurons will produce another burst of spikes. This is the “action” command that generates the eye movement. Most neurons that produce a movement-related burst also exhibit a sensory (visual) burst and thus are called visuomotor neurons. The sensory and motor bursts are temporally separated from each other in the delay task but can occur in close temporal proximity and even overlap partially in the gap task (e.g., Munoz and Wurtz, 1995; Sparks et al., 2000). Our goal was to determine whether the population SC activity in the sensory and motor response epochs are task dependent.

Previous neurophysiological studies of SC have mainly used single electrodes. Compared to the delay task, the gap task can evoke a slightly stronger sensory response for high luminance stimuli, especially when the target location and onset time are highly predictable (Marino et al., 2012). The motor burst, in contrast, is essentially comparable across tasks, provided that the saccade is directed to a visible target (Edelman and Goldberg, 2001). We applied machine learning-inspired dimensionality reduction techniques to identify the latent factors and subspaces spanned by the visual and motor responses of neural populations in both tasks, aiming to expand on the knowledge previously provided by single-electrode studies. Discriminant analyses were then used to determine the separability of the response distributions across epochs and tasks. The results verified that the visual responses were separable from motor responses in the delay task (Jagadisan and Gandhi, 2022), and a similar result was observed for the gap task. Interestingly, we also found across-task separability in both visual and movement responses. We interpret this latter finding to indicate that cognitive context can impact both exogenous and endogenous components of the sensory response as well as the transient features of the motor burst.

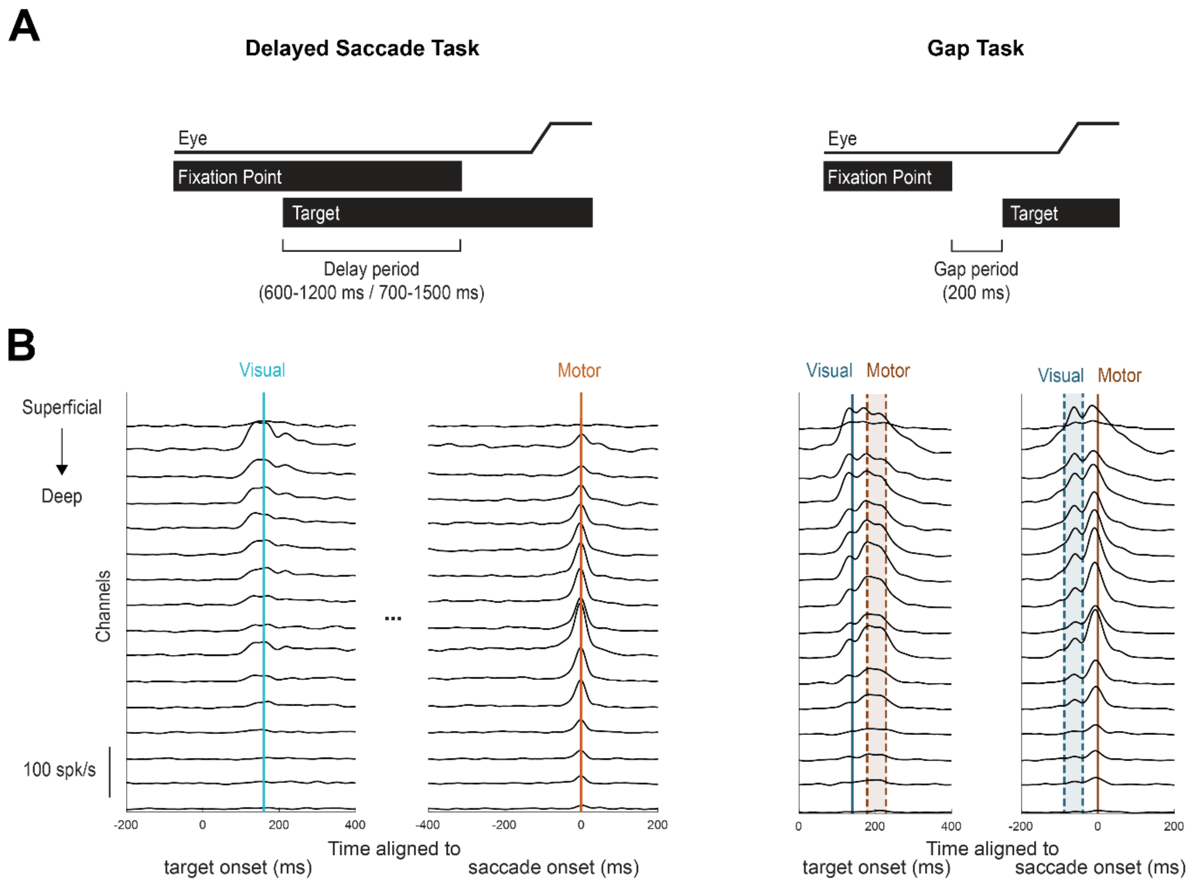

Figure 15. Overview of tasks and corresponding neural firing rate properties.

A. Flow of a single trial of the delay task (left) and gap task (right). The two tasks have identical key events (fixation on a central point, subsequent target appearance, and instructed saccade) that occur with different timing in each task. B. Multiunit activity averaged across trials for each recorded channel of one example session arranged in order from superficial to deep along the dorsoventral axis of the SC. The activity on each channel is aligned to target onset (left panel of each task) and saccade onset (right panel of each task). Activity around the peak of the visual and motor bursts are indicated by blue and orange vertical lines, respectively. For the gap task, the shaded rectangle in each panel corresponds to ±1 standard deviation around the mean saccade onset time (left panel) or visual burst time (right panel).

3.4 Methods

3.4.1 Animal Preparation

All procedures were approved by the Institutional Animal Care and Use Committee at the University of Pittsburgh and followed the guidelines of the Public Health Service Policy on Humane Care and Use of Laboratory Animals. Data were collected from two adult male Rhesus monkeys (Macaca mulatta, identified as BL and SU) who were trained to perform oculomotor tasks using operant conditioning and surgically prepared under aseptic conditions for neurophysiological recordings from the SC. The recording chamber was tilted 40 degrees posterior with respect to vertical to allow microelectrode penetrations orthogonal to the SC surface. A subset of the data used in this manuscript are a new analysis of data reported previously, and additional methodological details can be found there (Massot et al., 2019).

3.4.2 Experimental Design and Statistical Analyses

Stimulus display and acquisition of behavioral data were controlled by a central, customized program written in LabView architecture (Bryant & Gandhi, 2005). Eye position was measured using a camera-based pupil tracker (EyeLink 1000, SR Research, Ltd.) sampled at 1 kHz. Each animal was trained to perform two oculomotor tasks. Every trial was initiated by the presentation of a fixation point at the center of the screen, and the animal was required to direct its line of sight on this stimulus. In the delay task (Figure 15A, left panel), a single target was presented at an eccentric location while the monkey continued to fixate the central stimulus. The fixation point was extinguished after a “delay” period (randomized between 600-1200 ms for monkey BL and 700-1500 ms for monkey SU), permitting the animal to bring its visual axis within two degrees of the target. In the gap task (Figure 15A, right panel), the fixation point was first turned off, and a constant 200 ms gap in time lapsed before a single target was presented on the screen. The animal could not break fixation during the gap period but could make the saccade immediately after target onset. In both tasks, the animal received a liquid reward for maintaining fixation within a 2-3° window around the final target location for at least 250 ms. Trials of each task were either randomly interleaved (6 total included sessions across both monkeys) or presented in blocks (14 sessions across both monkeys).

Neural recordings in the SC were performed with a 16- or 24-channel linear microelectrode array (AlphaOmega Inc., or Plexon, Inc., respectively) lowered hydraulically in the SC to record neural activity. The probe trajectory was approximately orthogonal to the SC and thus traversed its dorsoventral axis. Neural activity was recorded via the Grapevine Scout Neural Interface Processor (Ripple, Inc.), visualized with associated Trellis software, and communicated with our central data acquisition system. Neural activity was band-pass filtered between 250 Hz and 5 kHz to record spiking activity. Spike times were determined using a standard threshold. Spike trains and spike density waveforms were inspected online for task-related activity characteristic of the SC when a target was presented near the center of the response field. We'll refer to the activity recorded across these contacts as each session's neural population. We also interleaved trials in which the target was placed at the diametrically opposite location, but those data are not considered here.

All analyses were performed using custom code written in MATLAB (MathWorks, Inc.). Saccades were detected using a 30-50°/s velocity criterion. For visualization purposes, we convolved the spike train with a Gaussian kernel of 10ms width to yield spike density waveforms.

For each session, we analyzed offline each channel’s waveforms on a trial-by-trial and trial-averaged basis to ensure each analyzed session consists of distinctive collicular activity across most channels for both the delay (Massot et al., 2019; Jagadisan and Gandhi, 2022) and gap tasks. To ensure that our results were not influenced by noise arising from too few trials, we set a criterion in which each session needed to have at least 30 trials from each task to be included. In most cases, each task condition had around 100 associated trials. Twenty (of 28) recording sessions met the above criteria (13 from monkey SU and 7 from BL) and comprise the dataset presented in this paper.

For each session, spiking activity from each channel was sorted into single units using MKsort, a supervised spike sorting interface (Ripple, Inc.). We found that both multiunit and single unit populations resulted in similar population activity patterns after dimensionality reduction, which is in line with a recent study (Trautmann et al., 2019). The results presented in Figure 2 through Figure 5 were obtained from spike-sorted neural populations.

A paired Wilcoxon signed rank test was performed on the two distributions of classification accuracy values in Figure 17 and Figure 18 to determine if the median difference between the two paired distributions was different from zero at an $\alpha = 0.05$ significance level. Additionally, a one-tailed Wilcoxon signed rank test was performed on each individual distribution in Figure 18 through Figure 20 to determine if the median classification accuracy was significantly above our pre-defined chance level of 50 percent. Last, for each individual session we performed an unpaired t-test on the distribution of values from the “delay” and “gap” conditions, as depicted in Figure 20A (saccade peak velocity) and B (saccade amplitude), followed by a paired t-test on the distribution of trial-averaged values from each condition across all sessions. See “Linear

Discriminant Analysis classification" subsection for details about the distributions on which these statistical tests were performed.

3.4.3 Definition of Sensory- and Motor-Related Activity

For both tasks, we wanted to compare neural activity across the sensation and action periods in two task contexts. Therefore, we describe sensation and action epochs as the respective times at which the neural population responds most strongly to a visual stimulus and contributes to movement generation in each task. For the delay task, we found through visual inspection of target onset-aligned, trial-averaged activity that the visual burst peak time occurred approximately 160 ms after target appearance for most neurons and sessions. Thus, we used this static time for all analyses of the visual epoch. The motor burst typically peaked around the time of saccade onset; thus, we used saccade onset time as the epoch from which we extracted motor activity. For the gap task, we found that the visual burst peak tended to occur slightly earlier, at approximately 140 ms following stimulus onset, and thus was used as the visual epoch for this task. As in the delay task, motor activity was taken from saccade onset in each gap task trial. Sensory-related activity is represented with (light/dark) orange and motor-related activity with (light/dark) blue (e.g., Figure 15B). These colors are used throughout the rest of our analyses comparing the neural activity underlying these key events.

3.4.4 Dimensionality Reduction Methods

Dimensionality reduction was used to summarize the population activity of simultaneously recorded neurons. This approach also offers better visualization of neural activity and comparison across tasks (Cunningham & Yu, 2014). Specifically, Gaussian Process Factor Analysis (GPFA) was performed on multi-channel spike trains aligned to target onset in discrete 20 ms bins of summed spike counts (Yu et al., 2009). Operating on a high-dimensional dataset, GPFA extracts a reduced number of factors that account for a large proportion of the variance and, additionally, smooths these factors in time. Temporal smoothing while simultaneously performing dimensionality reduction provides a more intuitive visual representation of how activity evolves over time. We applied the GPFA algorithm from DataHigh, a publicly available MATLAB code (Cowley et al., 2013), to a single structure, D, that pooled data from both tasks. For the delay task, we treated activity from 200 ms before target onset to 200 ms after saccade onset of each trial as one observation. For the gap task, we used activity from 200 ms before target onset to 600 ms after target onset (which includes the period of time in which a saccade is made) as one observation. D consists of an UxT matrix where U represents the number of units, or neurons, and T represents the number of 20ms time bins. After inputting D into DataHigh, a single low-dimensional representation is acquired that consists of both tasks' latent activity in a new reference frame – dimensions for which the highest variance in population neural activity across trials is found. Thus, our structure is converted into having a $U'xT$ matrix for each trial, where U' now represents the number of latent dimensions retained. We performed cross validation on our dimensionality-reduced data and found that the top 2-3 latent dimensions captured a substantial amount of the variance; therefore, we chose U' to equal 3 for every session for consistency. To note, since spike counts were binned in discrete 20 ms windows aligned to target onset rather than saccade onset, we aligned the binned latent activity in Figure 16A (right panels of each task) to the 20 ms bin in which the saccade was initiated on that trial for visualization, even though saccade onset could occur at any time in that window. Therefore, the latent activity appears to reach its trough slightly later than saccade onset; however, this should have no bearing on any across-epoch or across-task comparisons of neural activity, as all activity was processed in the same manner.

When initially plotting state-space activity for each session, we observed that a few sessions ($N = 3$) had a puzzling subspace "echo" within either the delay or gap task, in which activity from a subset of trials was shifted in one or more latent factors. When plotting the latent activity values against trial number, we saw for these sessions an abrupt trial on which the subspace shifted, potentially indicative of an electrode shift or the addition/loss of signal on one channel. For these sessions, we retained only the subset of trials for which the delay and gap task subspaces remained constant, discarding trials before or after the echo.

When performing GPFA on delay task trials only and then projecting gap task trials into the same low-dimensional space, we found qualitatively similar subspaces. Observations resulting from this analysis are not shown. In this case, we had two separate structures, one for the delay saccade task (D_1) and one for the gap task (D_2). GPFA was performed on the delay task first and then the projection weights calculated from this first (and only) occurrence of GPFA were retained so that gap task activity could be projected into the same low-dimensional space as the delay task. We also compared subspaces using matched visual burst times for both tasks (140 ms after target onset) and found them to be qualitatively similar to those formed by activity at the estimated visual burst peak time for each task (160 ms for delay; 140 ms for gap), indicating that our analyses of population activity patterns are not highly sensitive to the chosen event time during the visual epoch.

Supplementally, factor analysis (FA) was also performed using varying Gaussian smoothing kernels (width = 0, 6, and 10 ms) to ensure that the inherent temporal smoothing component of GPFA did not artificially constrain the separability between visual and motor activity in the gap task when the activity patterns may diverge at a short temporal scale. We found that manually choosing a smoothing parameter of 6ms before FA yielded qualitatively similar low-dimensional subspaces as those produced by GPFA. When performing FA with a longer kernel width of 10ms, the visual and motor subspaces become even more separable (observations not shown). This control analysis indicates that our results do not hinge on the exact decoding algorithm and parameters used.

3.4.5 Linear Discriminant Analysis Classification

A two-class linear classifier (linear discriminant analysis, LDA) was implemented to determine the amount of separability between various pairs of subspaces in the low-dimensional state space found via GPFA. We trained the classifier using a random partition of 70% of the data and the assigned class labels (“delay” and “gap” or “visual” and “motor”). The classifier was then tested on the remaining data points. LDA defines a decision boundary, or hyperplane, that maximizes the separability between the two classes. Classification accuracies were found by determining the number of instances in which the classifier accurately identified a point associated with a given label and then dividing by the total number of data points that should have been classified as such. To confirm statistical significance of these results, we performed 100 iterations of LDA on each data session. Presented in the results are the average classification accuracies taken across the 100 iterations for each data session. In addition, a control analysis was performed in which the categorical labels were randomly shuffled before the classifier was implemented to confirm our chance level is near the expected 50%. The orthogonality between the LDA planes separating visual and motor activity and delay and gap activity was determined by finding the angle between the two vectors normal to each plane. This angle was calculated by taking the arccosine of the normalized dot product of these vectors.

3.5 Results

In this study, we set out to characterize the patterns of population activity in the SC in two commonly used behavioral tasks to better understand contextual effects on neural signatures of sensation and action. In each task, both a sensory (visual) and motor burst is evident in the activity of SC neurons as the animal perceives and subsequently makes a saccade to a target that appears in his periphery. The structure of the delay task imposes a clear temporal separation between the two epochs while the gap task creates a situation in which the two bursts can overlap or even merge into one (Figure 15B). To quantify whether activity is distinct across epochs and/or tasks, we first performed GPFA to summarize the population activity in a less noisy, better visualizable framework. In doing so, we created a low-dimensional representation where each latent factor is a weighted linear combination, or summary, of all individual neurons included in that session's population (as in Figure 16A). In all datasets, like this representative example, the first latent factor qualitatively resembles the firing rate profile (the polarity of the signal is not crucial). The second factor, in contrast, features opposite deflections in the visual and motor epochs, although this is more appreciable when the two periods are separated in time as in the delay task. The third and higher dimension factors show minimal or negligible modulation. For each session, the absolute projection weights of each neuron were gathered for each dimension and all neurons were sorted by their projection weight magnitudes (Figure 16B). For all dimensions and sessions, we find that many neurons contribute to the variance, indicating that any resulting separation in activity is not dominated by a small number of individual neurons but rather due to the combined activity of a neural population.

The top 4 latent factors accounted for at least 95% percent of the variance in the firing rate profiles for 18/20 datasets. However, we chose to analyze only the top 3 latent dimensions because the third and higher dimensions do not show much modulation during the visual and motor epochs. This low dimensionality facilitates inspection and comparison of the factors across different conditions in a three-dimensional state space. Figure 16C illustrates in schematic form the "subspaces" spanned by the latent factors in the visual and motor epochs for the delay and gap tasks. Following this step, we implemented a two-class linear classifier to quantify the separability between each pair of neural subspaces. The classifier quantified the separability of activity 1) across epochs within each task (i.e., visual vs motor activity), 2) across tasks within each epoch (i.e., delay vs gap activity), 3) across epochs when both tasks are combined, and 4) across tasks when both epochs combined.

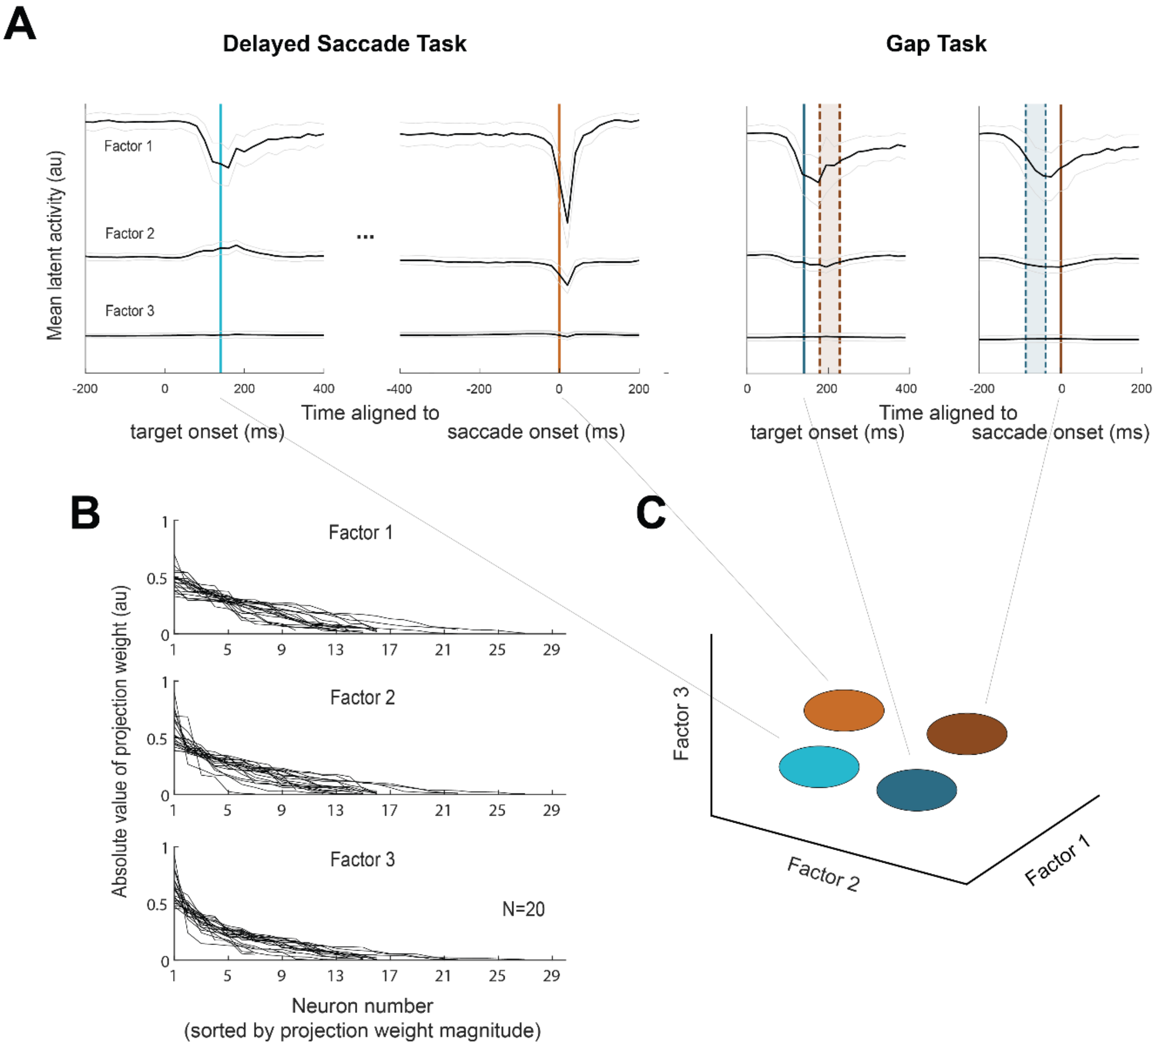

Figure 16. Overview of dimensionality reduction to view population activity in a state space framework. A. Magnitude of trial-averaged latent activity (± 1 standard deviation) across time for each of the first three factors retained after performing Gaussian Process Factor Analysis (GPFA) on spike-sorted population activity. We used the same conventions and example session here as described in Figure 15B. B. Projection weight of all neurons into each latent factor sorted from the neuron with the largest projection weight to that with the smallest for each session. Each session contributes one trace to each factor (N = 20). Most sessions have many neurons contributing to each latent factor; therefore, the three-dimensional latent values produce a true population-level framework rather than one dominated by a select few neurons. C. Hypothetical subspaces formed by activity during the visual (blue) and motor (orange) epochs of the delay task (lighter colors) and gap task (darker colors). The patterns produced across epochs and/or across tasks could be

nearly identical, producing overlapping subspaces, or they could be unique, resulting in separable subspaces as schematized here.

3.5.1 Within-Task, Across-Epoch Analysis

Determining if patterns are distinct across epochs for each task may give intuition as to whether there are unique neural signatures during two key points of sensorimotor transformation – sensation and action. As expected based on previous work (Jagadisan and Gandhi, 2022), the visual subspace is separable from the motor subspace for the delay task (light blue and light orange circles, respectively, in Figure 17A, left panel). We now report that the visual and motor subspaces are also distinct in the gap task (dark blue and dark orange diamonds, respectively, in Figure 17A, right panel). This separability likely reflects the varying contributions of each neuron during each epoch, as some SC neurons have larger visual bursts than motor bursts (and vice versa). Overall, regardless of the presence (or lack thereof) of an imposed delay, this session's population activity patterns are distinct between the sensory and motor epochs.

We performed LDA to quantify the separability between these two subspaces for each session (see Methods). The separability across epochs within each task was determined via a two-class linear discriminant classifier (with "visual" and "motor" categories). The classifier performed significantly above chance level across all sessions (one-tailed Wilcoxon signed rank test, $p < 0.001$), with a mean classification accuracy of 93.2% (±2.19%) for the delay task and 85.3% (±2.87%) for the gap task (Figure 17B). The classifier was significantly more accurate in distinguishing between sensory- and motor-related activity within the delay task than within the gap task (Wilcoxon signed rank test, $p < 0.001$). These results suggest that the between-epoch subspaces are more separable within the delay task, likely due to the imposed delay period which clearly separates the visual and motor epochs, as supported by previous research (Jagadisan and Gandhi, 2022). Overall, the separability of visual and motor patterns within a specific behavioral context is characteristic of a unique neural population activity pattern during the respective epochs.

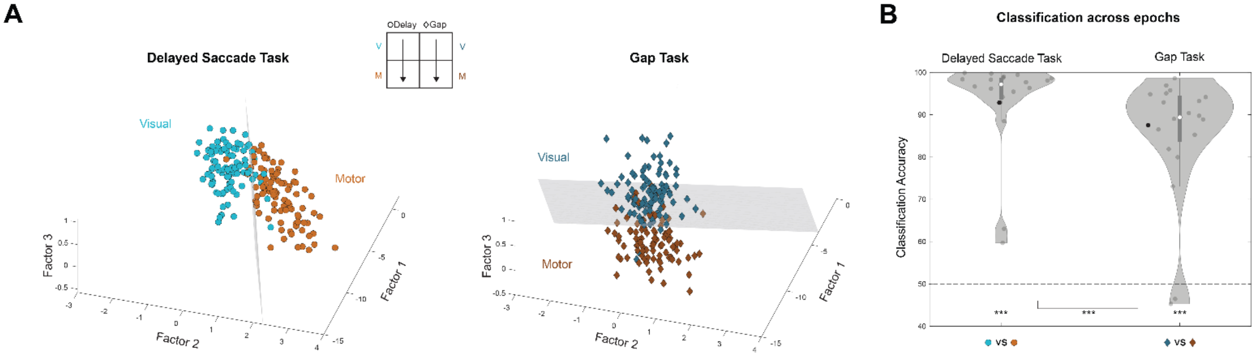

Figure 17. Separability of visual and motor subspaces within a task condition.

A. Clusters or subspaces of latent activity in the three-dimensional state space for an example session taken from the visual (light/dark blue) and motor (light/dark orange) epochs as defined in Figure 15B. Each point represents the latent activity in a 20ms window taken from an individual trial of the delay task (left panel) or gap task (right panel). The plane of maximum separability between the two clusters found through a linear discriminant analysis model (LDA) is denoted by a gray rectangle. The separability between the visual and motor subspaces was determined independently for each task. B. Violin plot of classification accuracy between the visual and motor epochs for the delay (left) and gap (right) tasks for all 20 sessions. Individual sessions are represented by gray points, with the example session indicated by a black point. A traditional box and whiskers plot is overlayed on the vertical meridian. Significance is reported both within and across distributions (one-tailed Wilcoxon signed rank test against a median of 50% and two-tailed paired Wilcoxon signed rank test, respectively). Accuracy in distinguishing visual from motor latent activity is significantly above chance level of 50% for both tasks (*p* < 0.001).

3.5.2 Within-Epoch, Across-Task Analysis

This analysis was then extended to compare patterns of activity between the two tasks within a given epoch. By comparing activity across tasks, we can gain a better idea if cognitive context is uniquely encoded in the superior colliculus. In Figure 18A, one can see that for the same example session as in Figure 3, there is indeed a separation during the visual epoch between the patterns exhibited by this neural population in the gap and delay tasks (blue circles and diamonds, respectively). A similar effect was also found for the motor epoch (orange circles and diamonds). Again, employing a two-class linear discriminant classifier (with "delay" and "gap" categories) to quantify the subspace separability for all sessions, we found the classification accuracy to be significantly above chance level during both epochs (one-tailed Wilcoxon signed rank test, $p < 0.001$). The across-session mean classification accuracy for the visual epoch was 77.9% (±3.82%) and was 80.8% (±3.36%) for the motor epoch (Figure 18B). There was no significant difference in the between-task separability of visual-related or motor-related activity (Wilcoxon signed rank test, $p = 0.24$). The presence of two distinct subspaces across tasks suggests that the way in which content-related signals are processed in the SC are unique to the cognitive demands of a given task.

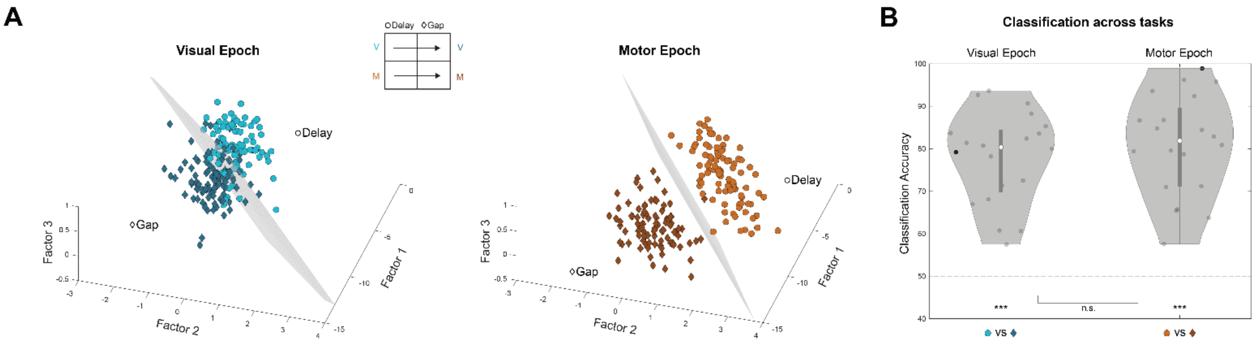

Figure 18. Separability of task-specific representations within an epoch.

A. Same conventions as in Figure 17A, but with planes of maximum separability computed between activity from delay (circles) and gap (diamonds) trials for the visual (left panel) and motor (right panel) epochs. B. Same as in Figure 17B, but for accuracy in classifying task condition during the visual (left) and motor (right) epochs. Accuracy in distinguishing activity from the delay task from that during the gap task is significantly above chance level of 50% for both epochs ($p < 0.001$).

3.5.3 Combined Across-Task and Across-Epoch Analyses

Lastly, we combined each task's activity and compared across epochs to determine if all visual activity is distinct from all motor activity regardless of the task condition. Figure 19A shows for the same example session the two-dimensional plane that best separates all visual activity from all motor activity. Across all sessions, there are distinct patterns of activity between the visual and motor epochs with a mean classification accuracy of 87.5% (±1.98%), which is significantly above chance level (one-tailed Wilcoxon signed rank test, $p < 0.001$) (Figure 19C). We also found that a single plane can considerably divide delay and gap task activity regardless of the epoch from which the activity came (Figure 19B). This classifier also performed significantly above chance level across sessions (Wilcoxon signed rank test, $p < 0.001$), with a mean accuracy of 72.7% (±2.85%) (Figure 19E).

In order to determine if SC populations encode content (visual vs. motor) and context (delay vs. gap) independently, we computed for each population the angle between the two vectors normal to each LDA plane. Figure 19D shows the distribution of these angles across all sessions. The median angle between decoding axes was 55.7 degrees, indicating a distinct but not orthogonal representation of content and context in these SC populations.

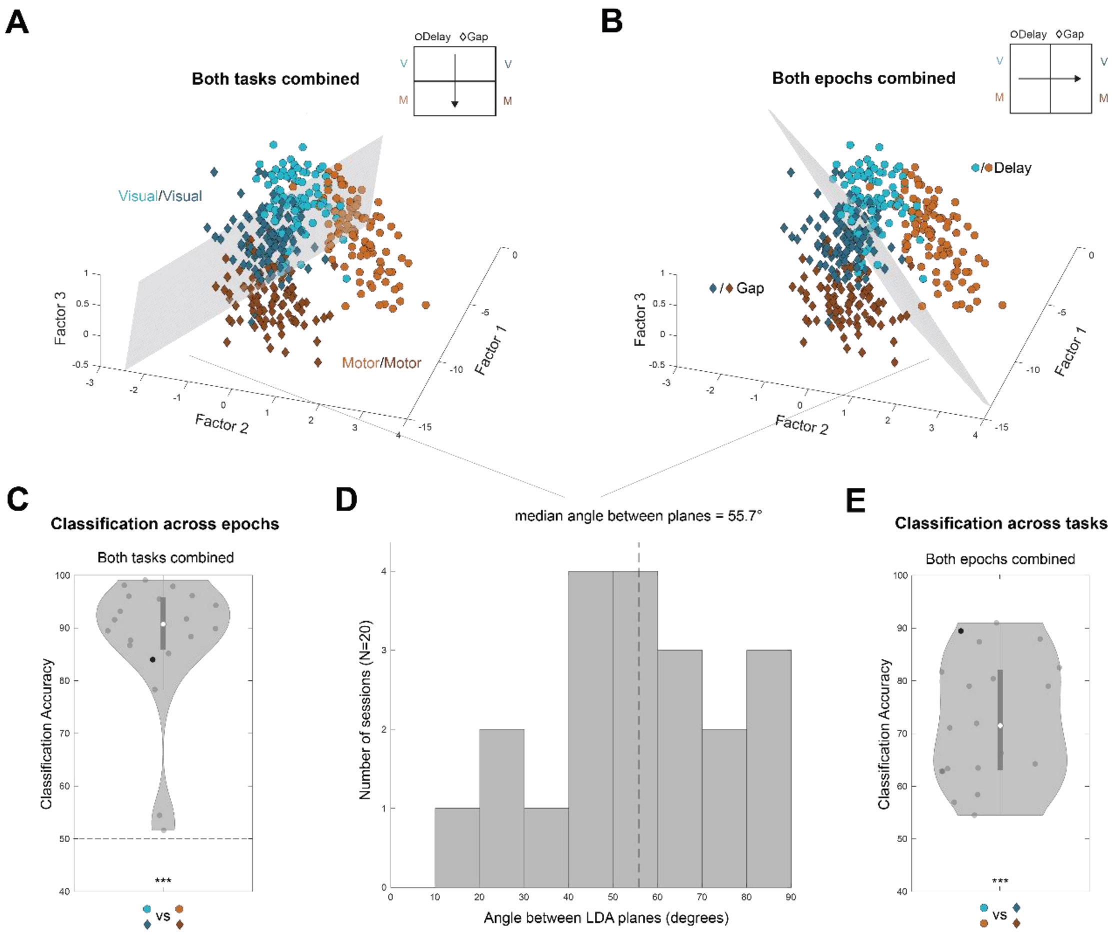

Figure 19. Combined separability of subspaces across epochs or tasks.

A and B. Same as in Figure 17A, but for (A) the separability of activity between the visual (blue) and motor (orange) epochs when activity is pooled across both tasks, or (B) the separability of activity between delay (circles) and gap (diamonds) tasks when activity is pooled across both epochs. C and E. Same as in Figure 17B, but for (C) accuracy in classifying visual and motor activity when both tasks are combined or (E) accuracy in classifying task condition when both epochs are combined. In both cases, the separability is significantly above chance level of 50% (*$p < 0.001$*). D. Histogram across sessions of the angle between two planes that maximally separate the content (e.g., gray plane in A) and context (e.g., gray plane in B) subspaces.

3.5.4 Behavioral Metrics

To ensure that the separability in the motor subspaces across tasks is not primarily attributable to a mismatch in behavioral metrics between the delay and gap tasks, we performed an additional set of analyses comparing the saccade peak velocities and saccade amplitudes across the two task conditions. Across sessions, the mean peak velocities for each task condition were significantly different (paired t-test, $p < 0.05$) (Figure 20A), but the difference was small and not likely of physiological significance. We think it is unlikely the cause of subspace separability across tasks given that separable subspaces were also observed in sessions without a significant difference in peak velocity. Saccade amplitude was not statistically different between the two conditions (paired t-test, $p > 0.05$) (Figure 20B).

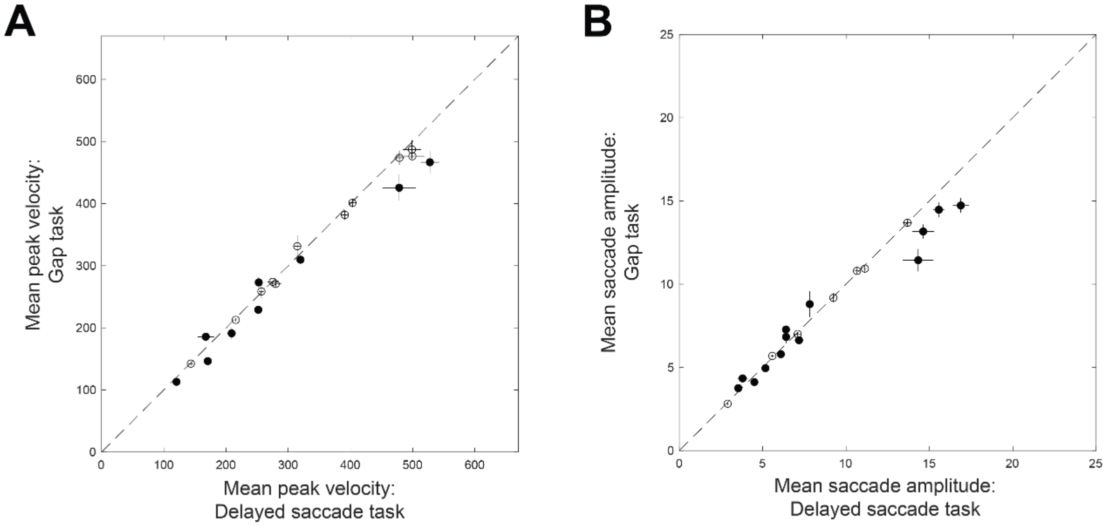

Figure 20. Comparison of behavioral metrics across task conditions.

A. Trial-averaged saccade peak velocities with 95% confidence intervals for the delay (x axis) and gap (y axis) tasks. Each session contributes one point, and the circle is filled in if the saccade peak velocity distributions between the two tasks were significantly different for that session (*p* < 0.05, t-test). Across sessions, the difference between the mean peak velocities in the two task conditions did reach significance (*p* = 0.038, paired t-test). B. Same as in (A) but for trial-averaged saccade amplitudes across tasks. Across sessions, the difference between the mean saccade amplitudes across tasks was not significant (*p* = 0.167, paired t-test).

3.6 Discussion

The superior colliculus is a central hub in the neuraxis that transforms sensation into action while incorporating cognition (Gandhi and Katnani, 2011; Krauzlis et al., 2013; Das, 2016; Basso and May, 2017; Cooper and McPeek, 2021). The general chronological order is that stimulus presentation induces a transient, exogenous burst of activity in SC neurons. Endogenous processes then contribute to SC activity, usually in the form of low-frequency modulation. This modulation has been linked to various cognitive features, as dictated by the experimental paradigm (D. L.

Sparks, 1999). Some examples include spatial attention (Lovejoy & Krauzlis, 2010), motor preparation (Everling et al., 1999; Jagadisan & Gandhi, 2017), anti-saccade generation (Dyckman et al., 2007), decision making (Crapse et al., 2018; Keller et al., 2005), visual search (McPeek and Keller, 2002; Shen et al., 2011), reward processing (Ikeda & Hikosaka, 2007), and remapping (Dunn et al., 2010). The culmination of the contextual processing leads to another burst of activity that produces a movement. Crucially, there is little consideration of context-dependent processing on the initial sensory and final motor bursts. Moreover, most studies have analyzed firing rates of individual SC neurons and then compiled the results by averaging the measures across neurons. The ability to record activity from many neurons simultaneously permits a higher order, dynamical systems analysis that examines covariability across neurons (Cunningham & Yu, 2014). Such efforts have yielded deeper insights into neural systems, although this technique has scarcely been applied to brain structures involved in producing visually-guided eye movements (Darlington & Lisberger, 2020).

In this study, we explored how neural populations in the SC encode two different forms of information – content (sensation/action) and context (two comparable behavioral tasks). First, we compared the population activity patterns during the periods of sensation and action and observed content-specific encoding in both the delay and gap tasks. This result matches closely with our intuition based on the known properties of SC neurons. Although we recorded primarily visuomotor neurons (i.e., those exhibiting both a transient visual and motor burst), the relative firing rate of each individual neuron during each epoch is imbalanced. Along the dorsoventral axis of the SC, neurons located more superficially respond more strongly to the appearance of a visual stimulus, while deeper neurons fire more strongly for the generation of a saccade (Massot et al., 2019). Since subspaces are formed by a weighted linear combination of activity across all recorded neurons, it logically follows that the population-level representations are distinct across epochs. This result is consistent with the notion that the trajectory formed by the neural population traverses through different regions of the state space as the trial transitions through stages of sensation, cognition, and action (Afshar et al., 2011; Churchland et al., 2010; Darlington & Lisberger, 2020) and is confined to lower-dimensional “manifolds” (Duncker & Sahani, 2021). We recently presented evidence for a temporal mechanism by which structures downstream of the SC could distinguish sensory from motor activity – even when the population-average firing rates are matched (Jagadisan and Gandhi, 2022). We also demonstrated that the temporal stability framework supplements the rate code by confining it to a subspace; thus our current and previous results align with each other. The present study therefore builds another case for the importance of characterizing neural activity through a measure that accounts for the relationship between neurons rather than by averaging the signal across all neurons, as is the standard in single unit studies.

Next, we explored the subspaces formed by population activity during the delay and gap tasks. During both the sensory period (equivalently, around the peak time of transient visual-related activity) and the motor period (around saccade onset), the subspace formed by the population activity pattern diverged across the two tasks, consistent with a context-specific representation. As the initial component of the sensory transient is considered exogenously driven, one may expect the visual subspaces to largely overlap for identical stimuli, but dimensionality reduction methods revealed a large separation for the two tasks (Figure 18). This is likely due to the interaction or addition of bottom-up visual transient with movement preparation-related activity that begins accumulating in the gap period, particularly when target locations and onset times are largely predictable (Dorris & Munoz, 1998; Edelman & Keller, 1996; Everling et al., 1999; Marino et al., 2022; D. Sparks et al., 2000), which is the case in our experiments. This accumulating activity can even be observed in early visual areas (Kim & Lee, 2017), which arguably serves to induce a different state in the sensorimotor circuitry by the time of target presentation and impacts the subsequent visual response in the gap task.

For matched behavioral metrics, the saccade-related motor burst of SC neurons is thought to be task invariant as long as the target remains illuminated (Dorris & Munoz, 1998; Edelman & Goldberg, 2001), although some exceptions have been reported (Hafed, 2021; Peel et al., 2021). Thus, we were surprised by the separation of motor subspaces across the two tasks (Figure 18). We were concerned that temporal smoothing component of GPFA, which prevents large and fast deviations of the neural trajectory, may have compromised occupancy of the actual motor subspace in the gap task because of the temporal proximity of the visual and motor bursts. Thus, we repeated the dimensionality reduction procedure with factor analysis (FA; see Methods) and found that the motor subspaces remained separated. We therefore believe the effect is physiologically meaningful, which can be appreciated by inspecting the spike density waveforms. In Figure 16B, for example, the 2nd and 3rd traces from the top show a stronger motor burst in the gap task. Such differences across neurons could account for the separation of motor subspaces across tasks. We speculate that this effect may arise because the animal updates his internal model of the relationship between neural activity and motor effector in response to the task context (e.g., Golub et al., 2015).

For our experimental design, we took considerable inspiration from a study by Lara et al. in which population activity patterns during reach preparation were found to be largely conserved across tasks with similar external events but that varied in both temporal and cognitive context (Lara et al., 2018). Here, we do not report the across-task neural patterns during the period in between sensation and action but rather during the sensation and action periods themselves.

Extrapolating our finding that the neural representations in the SC underlying sensation and action were not invariant to context, it would follow that activity patterns in the SC during the sensorimotor transformation process are also distinct across tasks. Additional studies that examine conversion of sensory into motor signals across multiple contexts are needed for oculomotor areas.

Since we found that SC activity is not only distinct in the sensory and motor epochs but also modulated by context, we lastly asked if these two dimensions of information are encoded not only simultaneously, but also independently. We found that within the scope of our study, these elements are not fully dissociable, as the planes that divide content 1 (visual burst) from content 2 (motor burst) and context 1 (delay task) from context 2 (gap task) are not orthogonal (Figure 19D). On the other hand, the planes are not parallel, which means that oculomotor areas to and from which the SC relays its signals might still decode task context somewhat separately from sensory and motor signals. In other words, knowing which task is being performed is not necessary for interpretation of epoch (i.e., deciphering whether the SC is signaling for target appearance or initiating a saccade command). Conversely, SC activity carries information about task context regardless of epoch. This dissociability of signals is analogous to a mechanism proposed by Kaufman and colleagues in which orthogonal encoding of preparatory and motor signals in dorsal premotor cortex prevents premature arm movement initiation (Kaufman et al., 2014). Hence, we believe that the application of dynamical systems approach is a valid and fruitful framework for future research into the representation of sensation and action in the oculomotor system across the multitude of possible behavioral contexts.

4.0 Decoding the Time Course of Spatial Information from Spiking and Local Field Potential Activities in the Superior Colliculus During a Sensorimotor Task

This work to be submitted as a manuscript with contributing authors: **Heusser MR**, Bourrelly C, Gandhi, NJ

4.1 Overview

How does the nervous system integrate sensory information into a goal-directed movement? This process, known as “sensorimotor transformation,” engages the superior colliculus (SC) for control of the eyes. Neurons across its dorsoventral axis exhibit transient bursts of activity following the appearance of a visual target and/or preceding a rapid eye movement to that target. The discriminability of spatial information encoded by neurons contained within a laminar column along this axis is thought to be limited to a narrow range of directions and amplitudes. We sought to challenge this notion by characterizing the time course of target direction tuning in two different signal modalities present in the SC – spiking activity and local field potentials. We recorded activity across many channels in a laminar column of the SC while rhesus monkeys (*Macaca mulatta*) performed a delayed saccade task to one of eight targets, which temporally separates the three main epochs of sensorimotor integration – visual, delay, and motor. To obtain a singular measure of spatial tuning across all channels, we employed a separate offline classification algorithm for each sliding time window and signal modality. For both spiking and LFP activity, decoding performance was highest during the visual and motor epochs. During the delay period, spiking activity exhibited spatial tuning akin to visual epoch tuning, while target direction was represented more broadly during the motor epoch. For LFP activity, spatial tuning was similar to that of spiking activity during the visual and motor epochs. Delay period tuning was largely absent in LFP signals, often confined to the target in the preferred direction. Overall, this study demonstrates that the encoding of direction information across layers of the SC is broad in tuning and dynamic in nature.

4.2 Introduction

We interact with our environment by redirecting our line of sight to objects of interest. A large network of neural structures is involved in this process of sensorimotor integration. The main neural pathway for producing fast eye movements, also known as saccades, is from the cortex to the midbrain, then to deeper nuclei in the brainstem, and finally to the end effectors (i.e., eye muscles). For any given saccade, populations of neurons are active and signal to varying degrees the spatial parameters (e.g., direction and amplitude) of the target and/or the intended movement, but the anatomical organization with respect to spatial parameters varies across areas. Placing multicontact electrodes into cortical oculomotor structures such as the frontal eye fields (FEF) yields a heterogenous population – each neuron in the recorded population will signal maximally for a very different amplitude and direction of the intended eye movement (Bruce et al., 1985; Sommer & Wurtz, 2000). The superior colliculus (SC) is a deep brain structure that receives direct input from cortical oculomotor areas and is thought to be a hub for sensorimotor integration (see reviews by Basso & May, 2017; Gandhi & Katnani, 2011). In contrast to cortical oculomotor areas, a typical electrode trajectory through the SC does not provide appreciable spatial variability.

Within a layer of the SC, neurons vary in their preferred saccade direction and amplitude along the mediolateral and rostral-caudal axes, respectively. Thus, the position on the SC topographic map of maximal activity across all active neurons (i.e., the “hot spot”) is what determines the direction and amplitude of the executed saccade. However, experimenters most often approach the SC with probes inserted orthogonally to the SC surface, with electrode contacts spanning the dorsoventral axis. Neurons along this axis systematically vary in the degree to which they signal for sensory or motor parameters across time and depth (Ikeda et al., 2015; Massot et al., 2019; Mohler & Wurtz, 1976) but are thought to encode roughly the same intended saccade vector and thus give a highly homogenous population in this regard.

Solving the inverse problem – decoding the intended saccade direction from activity recorded at one particular location on the SC’s topographic map – is not straightforward. Single unit studies would have us believe this is not possible; neurons located at a given point on the SC map fire maximally for a specific target or saccade direction, but any equal amount of change in direction from the preferred direction (say, ± 45 degrees of visual angle) results in roughly equal firing profiles and obscures any relationship with the external parameter. On the other hand, if we record from many neurons along the dorsoventral axis simultaneously, the pattern across the population may still contain information about spatial parameters across a broad range of the visual field despite the assumed homogeneity across layers.

Therefore, we set out to determine if information about saccade or target direction can be obtained from activities of small populations of SC neurons within a specific location on the SC topographic map, and if so, at what time(s) during the sensorimotor integration process this spatial information is present. Accordingly, we investigated the time course of direction tuning present in the spiking activity of SC neural populations and compared the neurons’ encoding properties with a second signal modality, the local field potential (LFP), which at a given recording site reflects the aggregate activity of nearby neurons through a measure of their extracellular voltage (Buzsáki et al., 2012). We recorded the spiking (Figure 21) and LFP (Figure 22) activity simultaneously across layers of the SC while rhesus monkeys (*Macaca mulatta*) performed delayed saccades to one of eight targets radially equidistant in direction. We then trained a simple linear classifier to output the most likely direction to which small windows of spiking or LFP activity belonged. The performance of the classifier across time and directions gives a comprehensive indication of the spatial encoding properties of SC activity during sensorimotor integration. We found that beginning in the early delay period, spatial information is more strongly encoded in spiking activity than in LFPs. For both modalities, the spatial tuning width was significantly broader in the motor epoch than in the visual epoch. Overall, these results indicate that broad directions are encoded within single columns of the SC and that the spatial encoding properties are distinct across signal modalities and epochs. Perhaps the inverse problem of broad spatial parameter decoding from the epicenter of SC activity is not a problem after all.

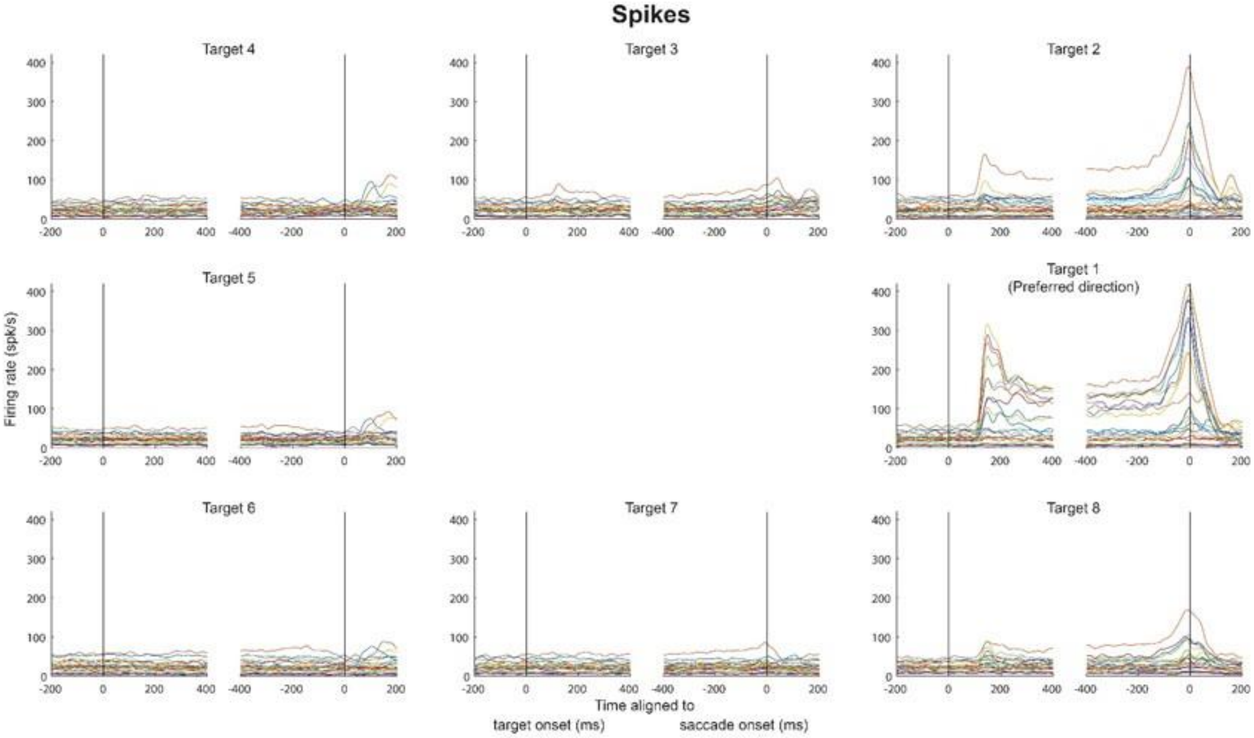

Figure 21. Peri-event time histograms of spiking activity simultaneously recorded across 24 channels. The across-trial mean firing rates for all 24 channels recorded during an example session are plotted aligned to target onset (left) and saccade onset (right) for trials to eight radially equidistant targets. Each colored trace represents the spiking activity on one channel averaged across all trials to a particular target. Subplots are rotated so that the preferred target direction of this population is displayed horizontal and rightward with respect to center.

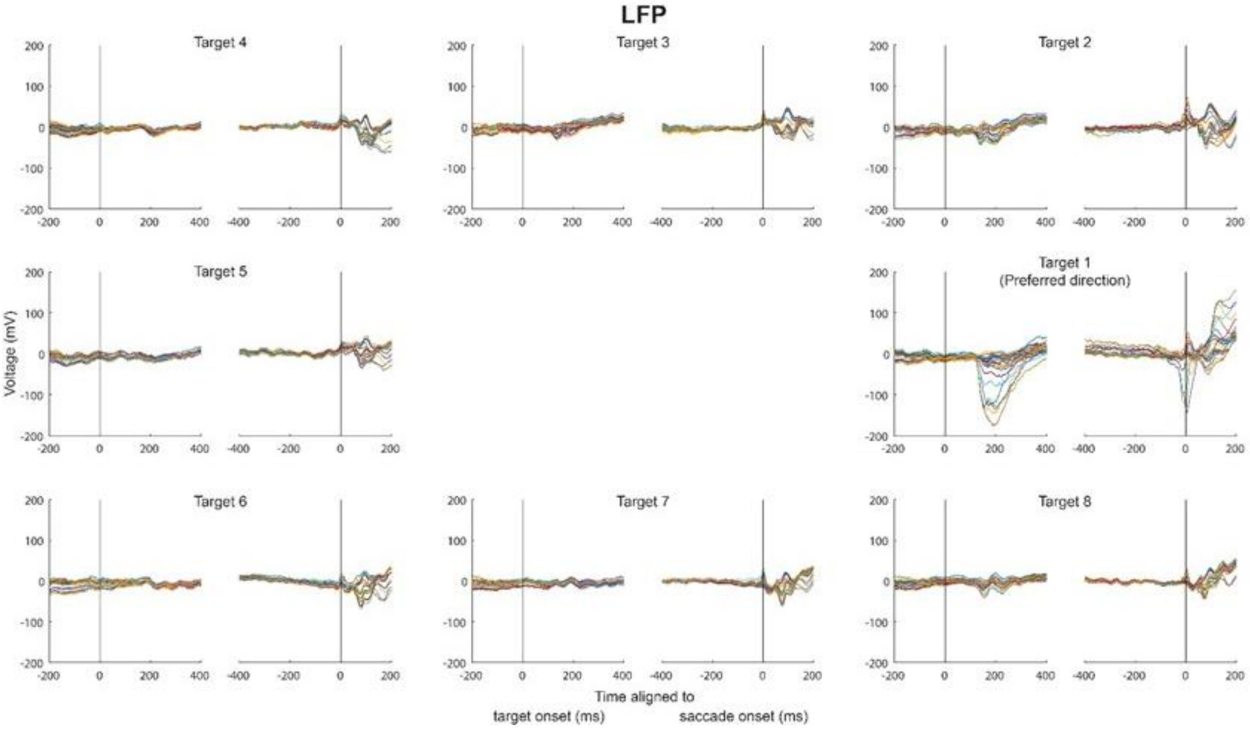

Figure 22. Peri-event voltage values of local field potentials (LFPs) simultaneously recorded across 24 channels.

The across-trial mean LFP voltage value for all 24 channels recorded during an example session are plotted aligned to target onset (left) and saccade onset (right) for trials to eight radially equidistant targets. All conventions are the same as in Figure 21.

4.3 Methods

4.3.1 Animal Preparation

Two adult male rhesus monkeys (*Macaca mulatta*; BL and SU) were used in this study. All experimental procedures were approved by the University of Pittsburgh Institutional Animal Care and Use Committee. A sterile surgery was performed on each animal to implant a stainless steel recording chamber (Narishige, Inc.) angled 40 degrees posterior with respect to vertical. Electrode penetrations through this chamber approach the SC orthogonal to its surface and traverses its dorsoventral axis along a track where neurons have similar response fields. Both animals were fitted with a thermoplastic mask to achieve fixation of the head during experimental sessions.

4.3.2 Data Collection

Comprehensive details about neurophysiology and microstimulation are provided in Massot et al., 2019. In brief, a 16 (monkey BL) or 24 (SU) channel laminar microelectrode (Alpha Omega, Inc., or Plexon, Inc., respectively) was inserted acutely into the SC to record neural activity across different layers. Many individual channels were stimulated to qualitatively determine an average evoked saccade vector, which was used as the preferred location for that session's neural population. The raw activity recorded on each channel was separated into spike times (high pass filtered at 250 Hz and discretized using a standard threshold) and LFP (low pass filtered at 250 Hz). For visualization only, spike counts were converted into firing rates by convolving each channel's spike train with a Gaussian kernel of 10ms width (as in Figure 21) and LFPs were bandpass filtered between 0.5 and 250 Hz with a notch filter at 60 Hz (as in Figure 22). The majority of channels with task-related spiking activity encountered visuomotor neurons that exhibited large transient bursts both in response to a visual stimulus and before/during saccade. Data from 15 sessions from monkey SU and 3 sessions from monkey BL were collected and all correct trials were included in analyses ($N = 18$ total sessions).

4.3.3 Behavioral Paradigm

Each monkey was trained to sit in a primate chair and perform a standard eye movement task in a dark room. Eye position was tracked with an infrared eye tracker (EyeLink 1000, SR Research, Ltd.; see Massot et al., 2019 for additional details). During each recording session, animals performed many trials of a delayed saccade task to one of eight possible targets evenly spaced in 45-degree increments around the fixation point. Each target had an equal likelihood of presentation, and “Target 1” was either placed at the spatial location corresponding to the estimated preferred saccade vector (for the majority of sessions) or at the position (10, 0) in Cartesian coordinates. In the latter case, preferred target direction was re-defined as Target 1 following examination of the average spiking activity profiles for that session (as in Figure 21). Temporal progression of a typical trial is shown in Figure 3. Important to note are the three main epochs that occur in each trial – visual, delay, and motor. The animal was given a liquid reward after successful trials, and only these trials were included in analyses (typically yielding over 1000 total trials across all target directions per session).

4.3.4 Classification Methods

Custom MATLAB code (MathWorks, Inc.) was used for all analyses unless otherwise specified. Target location was decoded offline from population activity on each session individually. Summed spike counts or average LFP voltage on each channel in 100 ms time windows (different lengths were also used, as described below) sliding across the duration of each individual trial (dt = 10 ms) were labeled as belonging to Target 1 through Target 8 depending on the target location presented on that trial. For each individual 100 ms time bin, a separate linear discriminant classifier was trained on these summed spike counts or average LFP voltage from a randomly selected 70% of total trials (pooled across all targets), and its performance was tested on the remaining 30%. Classifier performance was measured through the F1 score, a common metric for multiclass classifiers that takes into account both the sensitivity and precision of the model for each target, which counters any overfitting/underfitting of the model to activity belonging to particular target (e.g., Zhi et al., 2018). This process of randomly selecting 70% and 30% as training and test trials, respectively, was repeated for a total of 10 times for each window and each session to obtain an average classifier performance across iterations. Importantly, each classifier was trained and tested only on activity belonging to a particular time range and had no information about future or past windows that would influence performance within a given window. To determine an experimental chance level, target labels were randomly shuffled and the classification process described above was repeated. The actual chance level tended to closely match theoretical chance level performance of 1 out of 8 targets, or 12.5% (results not shown). Before averaging across sessions, the classifier performance value of true and shuffled data in each window for each target was subtracted by the mean performance value for that target during the first 200 ms of the baseline period (i.e., 400 ms to 200 ms before target onset). This was done to normalize all sessions' performance values as a change in performance relative to baseline. Therefore, in Figure 23 through Figure 27, performance values do not represent the absolute performance of the classifier (which often reached 1, the highest possible value).

A linear discriminant analysis (LDA) classifier is a supervised, geometric model that finds a hyperplane that maximally separates the input features between two categories, or "classes," during the training phase. In this paradigm, there are 24 input features that correspond to the spiking or LFP activity on each channel, but there are 8 classes that correspond to the 8 targets presented. Since an LDA model is by definition a binary classifier, we implemented a common technique called error-correcting output codes (ECOC) that fits a series of binary LDA classifiers in a one-vs-one manner to convert the model into a multiclass classifier, allowing for simultaneous classification into more than two categories (Derya Übeyli, 2008). During the testing phase, new data is shown to the model, and the class (i.e., target to which the activity corresponds) is determined by the position relative to the hyperplanes that were found during the training phase. To note, a pseudolinear discriminant classifier was implemented for spiking activity to combat the low or absent spike counts on some channels in certain time windows, which often leads to zero variance across observations and disrupts model fitting. We also repeated all analyses using a ECOC support vector machine (SVM) algorithm and found classifier performance dynamics for both spiking and LFP activity to be quite similar to results found via ECOC LDA classification.

For analysis of the effect of window length on classifier performance, spike counts were summed and LFP voltage was averaged across each window of length [20, 50, 100, 200, 300] ms, which again were calculated in sliding increments of 10 ms. In all visualizations of classifier performance across time (Figure 23 through Figure 26), values are plotted in a causal manner; for example, performance for the set of observations in the time window 100 to 200 ms after target onset is plotted at the 200 ms mark to represent that only historical activity was used to create and test a model of spatial location information.

To represent the spatial tuning properties of our neural populations during the many epochs of this behavioral task, we defined a range of times for each of five epochs (baseline, visual, early delay, late delay, and motor) during which we pulled out a single across-session classifier performance value for each target direction. Baseline performance was taken as the mean value in the range of 400 ms to 200 ms before target onset. Visual performance was taken as the maximum value around the time of the visual burst, typically occurring within the 100 ms to 200 ms range after target onset. Early delay performance was taken as the mean value in the range of 250 ms to 450 ms after target onset. Late delay performance was taken as the mean value in the range of 300 ms to 100 ms before saccade onset. Motor performance was taken as the maximum value around the time of the motor burst, typically occurring near saccade onset.

All statistical comparisons of classifier performance between signal modalities or epochs (in Figure 27 and Figure 28) were performed using a paired two-tailed t-test with $\alpha < 0.05$ indicating a significant difference between the two distributions included in the comparison.

4.4 Results

In this study, we set out to determine if and at what times during a behavioral task do neural populations in a single column of the SC encode information about saccade target direction. We employed a simple offline decoding algorithm (linear discriminant classifier) as a proxy for discriminability of spatial location (i.e., to which out of 8 possible targets will an animal make a saccade on a given trial) during independent sliding windows of time throughout a behavioral task.

This decoding algorithm was first applied to spiking activity of simultaneously recorded neurons and then to the local field potential (LFP) recorded at the same locations across the dorsoventral axis of the SC. Importantly, we remain agnostic with respect to whether the population encodes sensory and/or motor information at any given time. Instead, we will use any combination of terms “target/saccade direction/location” throughout the text and do not make any attempts to distinguish whether the spatial information being encoded is related to sensory (i.e., visual stimulus angle relative to eye position at fixation) or motor (i.e., intended saccade direction relative to starting eye position) representations.

Figure 23 shows the across-session mean performance in decoding target location from small windows of summed spike counts for each target. Here, Target 1 (middle right) has been rotated for each session to represent the target location most preferred by the neural population recorded on that day (as determined by microstimulation, see Methods). By aligning all sessions according to their preferred target location, we can better appreciate any change in decoding target location as a function of the proximity of a target to the preferred target. In other words, Targets 2 and 8 are approximately equidistant from the preferred target, while Target 5 represents the target diametrically opposite to the preferred target – one that is in the opposite hemifield.

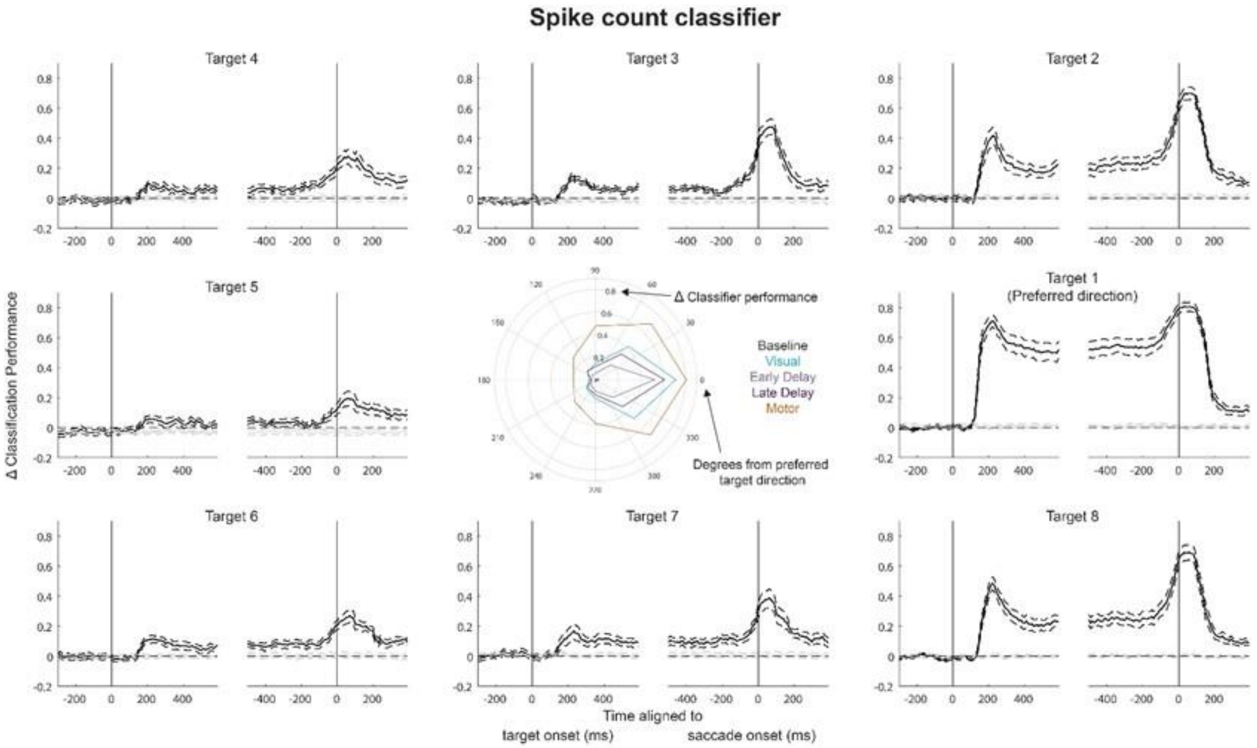

Figure 23. Linear discriminant classification of spiking activity.

Sliding 100 ms windows of summed spike counts on each channel were used to train a linear discriminant analysis (LDA) model and test its ability to decode target direction. Mean (±SEM) across-session classifier performance is plotted separately for each of eight target directions and aligned to target onset (left panels) or saccade onset (right panels). Chance level classifier performance was obtained by using shuffled class labels during the training phase. Performance values were grouped across sessions by aligning to each session's preferred target direction (visualized here as the right middle panels), and the performance for each session and each target was baseline-subtracted before averaging. Values for each window are plotted aligned to the end of that window (e.g., performance of the classifier trained and tested on the 0 ms to 100 ms window following target onset is plotted at 100 ms on the x axis). Inset: Summary polar plot of mean across-session classifier performance distribution across target directions during each epoch as defined in Methods. Spatial tuning of spiking activity is broader in the motor epoch than any other epoch.

The first, and perhaps most obvious, observation to note is that spatial information is best decoded during the neural populations' visual and motor bursts, peaking roughly 150 ms after target onset and again around saccade onset, respectively. This aligns well with the population-averaged response during these two epochs (Figure 21). Next, and perhaps just as intuitive, is the observation that the decoding performance is best for the target in the preferred location. Equivalently, the spiking activity pattern is most distinct from other target locations when the target is presented in the preferred location (i.e., the target that evokes a maximal firing rate in response to its appearance).

Another observation from Figure 23 is that spatial information can still be decoded from targets far away from the preferred location (e.g., Targets 4-6). Despite the low firing rate modulation for these targets, the spiking activity is in fact still distinct across targets presented in this region; otherwise, the performance would remain at baseline level (here, at 0 on the y-axis) throughout the trial. Instead, the classification performance is well above chance level for these directions, including for the location diametrically opposite the preferred direction. This result can likely be attributed to the activity seen in individual channels when targets in this region are presented, although the direction of modulation (i.e., elevation or suppression of activity) for saccade targets in this hemifield is unique to each individual neuron and population (see example session in Figure 21).

The last main observation in Figure 23 is that the decoding accuracy remains elevated throughout the delay period, in the time between the transient visual burst and the much-later motor burst, especially for targets in and near the preferred location. This result suggests that target location is one form of information still present during the delay period, which can be attributed to the sustained tonic activity exhibited by many SC neurons following the end of the transient visual response.

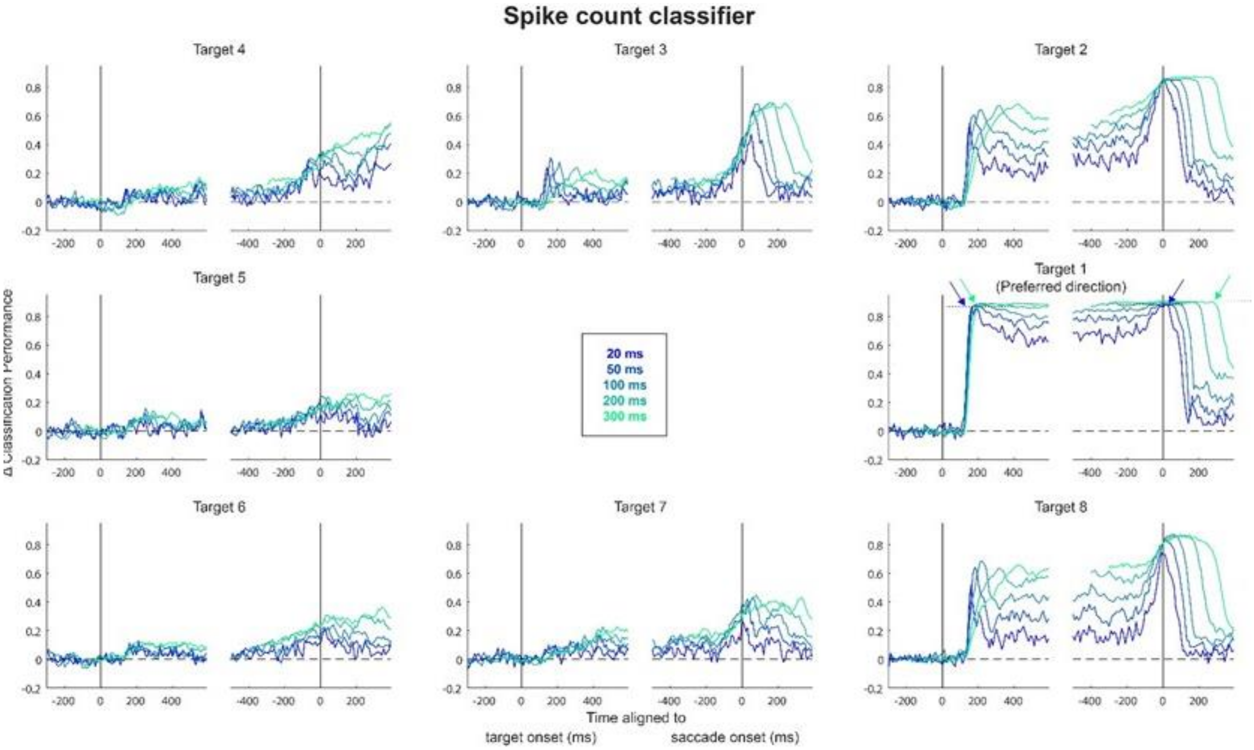

Figure 24. Linear discriminant classification of spiking activity: systematic variation of bin width. Classifiers were trained and tested on summed spike counts during windows of lengths ranging from 20 ms to 300 ms. Average performance over 50 bootstrapping iterations for each target direction and each window length condition are plotted using the same conventions as Figure 21 for one example session. Again, values are plotted aligned to the end of each window; therefore, each condition peaks in classification performance at different times but this is not the comparison of interest. Spike count-based classification is largely robust to window size during the transient visual and motor epochs (as indicated by the dark blue and green arrows at Target 1) but performance increases with increasing window sizes during the delay period.

We wanted to determine if these observations were robust to the size of the window used to classify the target location. Therefore, we systematically varied the bin width of summed spike counts used to train and test the classifier from very small (20 ms) to very large (300 ms), and the across-session mean performance for each bin width is shown in Figure 24 for one example session. Indeed, varying the bin width did not qualitatively change the conclusions drawn above.

Instead, the spatial location decoding performance gradually increased as bin width increased, plateauing around the 100 ms window length. In other words, using summed spiking activity from time ranges longer than 100 ms did not improve the classifier performance, from which we infer that information about spatial location is encoded maximally in short periods of spiking.

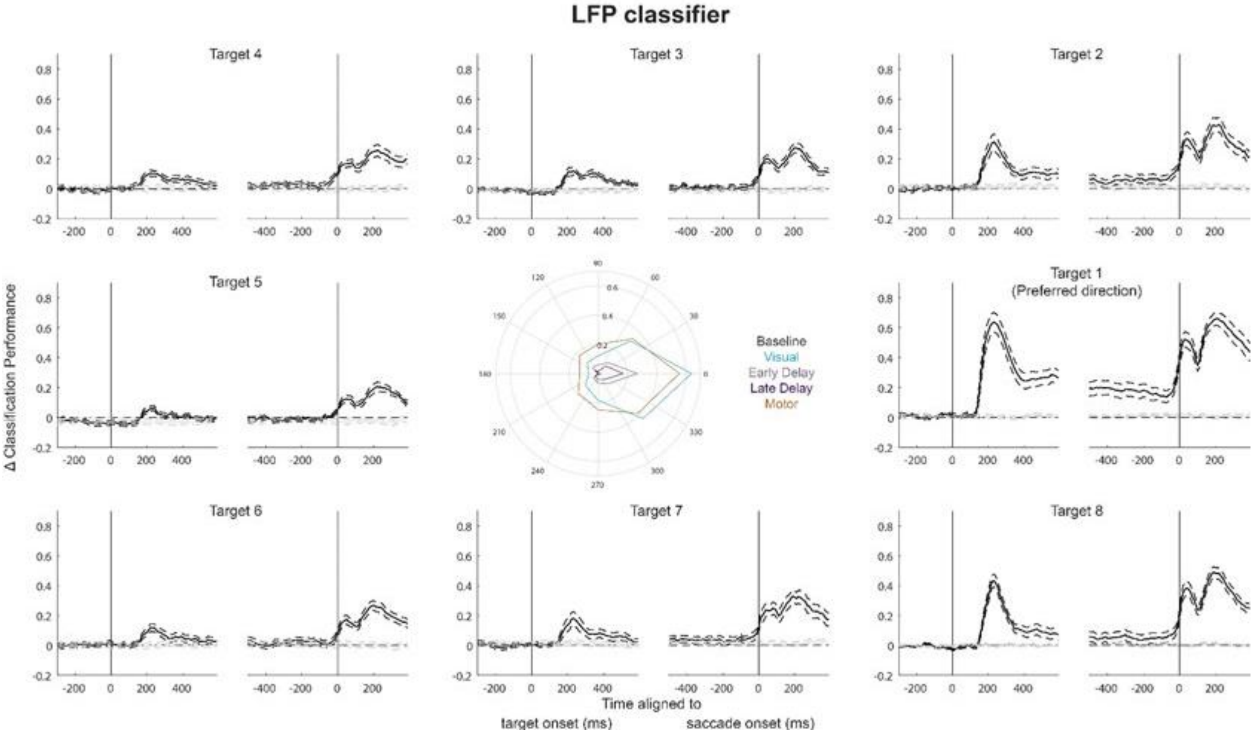

Figure 25. Linear discriminant classification of local field potentials.

Sliding 100 ms windows of average LFP voltage on each channel were used to train an LDA model and test its ability to decode target direction. Mean (±SEM) across-session classifier performance is plotted separately for each of eight target directions and aligned to target onset (left panels) or saccade onset (right panels). All other conventions are the same as in Figure 21. Inset polar plot shows that decoding accuracy is lower during the delay period but is consistent between the visual and motor epochs.

Next, we applied a classification algorithm to the local field potentials recorded simultaneously across many channels. Figure 25 shows the across-session mean performance when decoding target location from small windows of averaged LFP voltage signals. A decoding performance comparable to the spike count-based classifier was found during the visual epoch. However, in contrast to the spiking activity-based classification, the ability to decode spatial location from LFPs during the delay period is much more constrained to the preferred target direction. This tuning reverts to a broad one during the motor epoch, although the extent of spatial information does not expand past that observed during the visual epoch as it does in the spike-based classifier. When the LFP signal is averaged across windows ranging from 20 ms in length to 300 ms in length, as shown in Figure 26, we see that the maximum performance is reached when the window length is the shortest during the visual and motor epochs (see dark blue and light green arrows for Target 1). This short optimal window length suggests that spatial information is encoded maximally in short periods of time during these transient epochs, unlike that observed in the spike-based classifier. However, just as with spiking activity, spatial information seems to be maximally encoded on a longer time scale during the delay period.

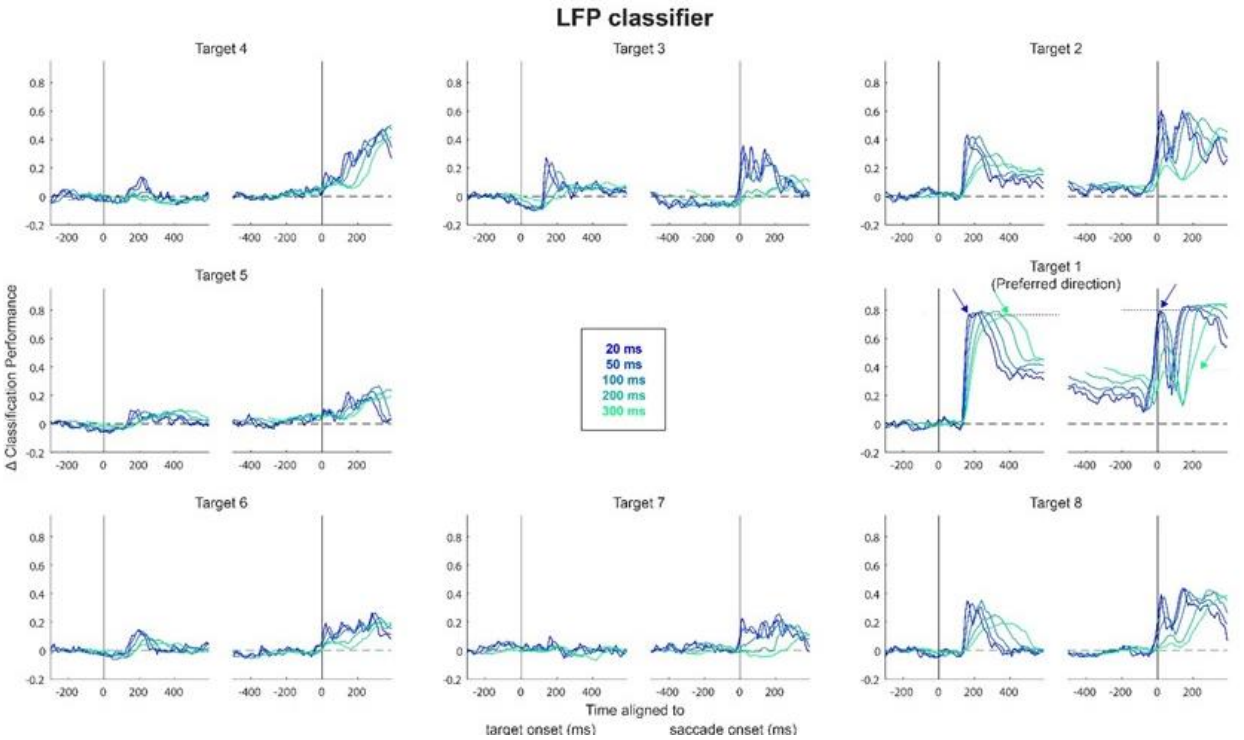

Figure 26. Linear discriminant classification of local field potentials: systematic variation of bin width. Classifiers were trained and tested on average LFP voltage on each channel during windows of lengths ranging from 20 ms to 300 ms. Average performance over 50 bootstrapping iterations for each target direction and each window length condition are plotted using the same conventions as Figure 21 for one example session. A decrease in performance with increasing window lengths can be seen during the motor epoch (indicated by dark blue and light green arrows at Target 1), but the opposite effect can be seen during the delay period.

Last, we quantitatively compared the spatial encoding properties across epochs and signal modalities – first for each individual target direction and then integrated across all eight target directions. Figure 27 breaks down the classification performance during the visual epoch vs. the motor epoch independently for each target and signal modality. The spike-based classifier produced consistently higher performance in the motor epoch than in the visual epoch for all target directions irrespective of the angular distance from the preferred location. On the contrary, the

LFP-based classifier only displayed significantly different performance between the visual and motor epochs for target directions far from the preferred direction.

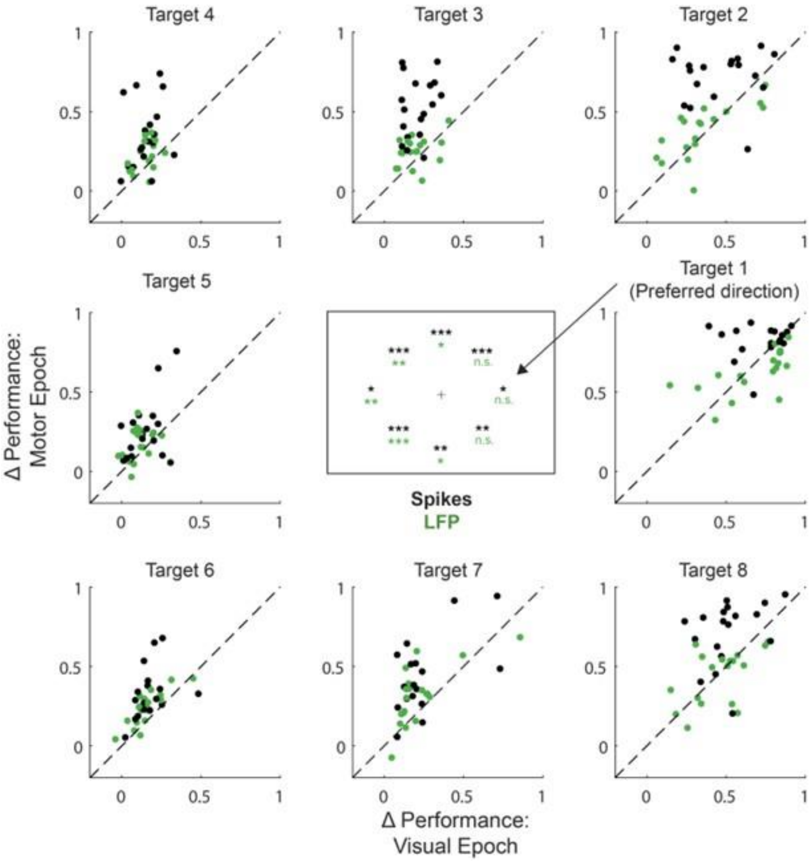

Figure 27. Comparison of direction encoding during the visual and motor epochs for each target. Average decoding performance in the visual (x-axis) vs. motor (y-axis) epoch as defined in Methods for each target. Spike-based classifiers are indicated in black and LFP-based classifiers are indicated in green. Each session (N=18) contributes two points to each of the eight target subplots – one for spiking activity and another for LFP activity. Inset: Significant (paired t-test) differences in performance level during the visual and motor epochs for each target are represented, with $p<0.05$ indicated by a single asterisk, $p<0.01$ by double asterisks, and $p<0.001$ by triple asterisks. For spike-based classifiers, the performance is significantly different between epochs for all targets. For LFP-based classifiers, only targets far from the preferred direction have significantly different encoding across epochs.

Figure 28A shows the integrated decoding accuracy across targets into four different epochs (see Methods for definitions) for the spike-based classifier in black and the LFP-based classifier in green. We computed the area under the tuning curves separately for each epoch and signal modality to obtain a proxy for overall accuracy, and these values are shown in Figure 28B. Beginning in the visual epoch, the amount of spatial information is significantly different between spikes and LFPs (paired t-test), and this separation persists throughout the time course of the trial. To obtain a measure of the classification tuning width – that is, the narrowness or breadth of ability to characterize activity across the full range of target directions – we shifted each population's decoding values such that the decoding performance was 1 for the target in the preferred direction (i.e., Target 1) before taking the area under the tuning curve. This provides a means of normalization across epochs so that any uniform shifts in decoding performance across all targets from one epoch to another do not impact this measure. The results in Figure 28C demonstrate that the spatial tuning width is only significantly different between signal modalities during the delay period (paired t-test). Comparing each signal modality's tuning properties across epochs (Figure 28D), we see that for the spike-based classifier, the tuning width is only significantly different between the visual and motor epochs and between the late delay and motor epochs. For the LFP-based classifier, the tuning width is significantly different across all epochs, indicating a dynamic shift in spatial encoding across epochs.

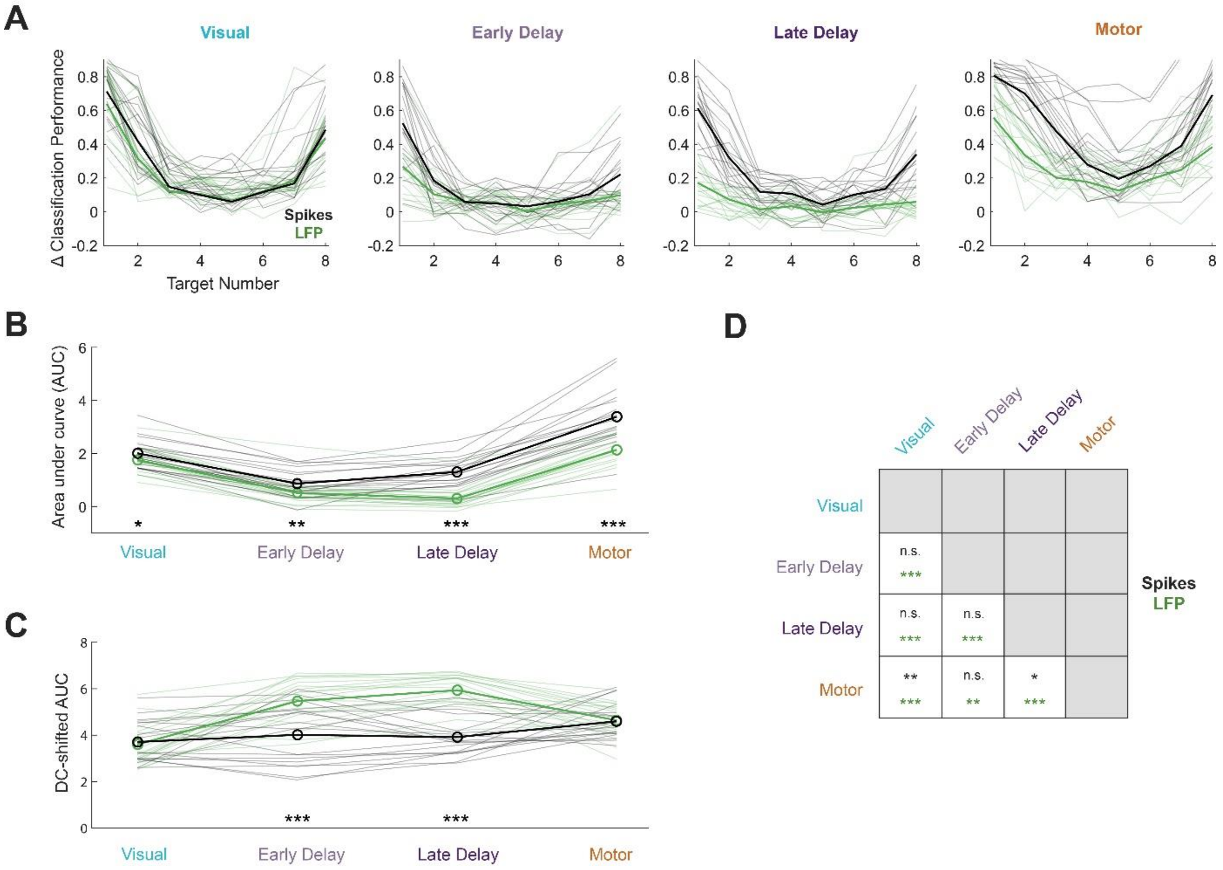

Figure 28. Comparison of spatial encoding properties of spiking and LFP activity across epochs. A. Baseline-shifted classification performance on spiking (black) and LFP (green) activity during each of the four main epochs (as defined in Methods) for each target aligned to the preferred direction of the population. Mean across sessions (bold lines) as well as each session's individual tuning curve (N=18, thin lines) are shown. B. Trapezoidal area under the observed tuning curve (AUC) during each epoch for each session (thin lines) along with the across-session mean AUC (bold lines). Significant differences between spiking and LFP classifier distributions using a paired t-test are shown with asterisks at the α=0.05 significance level (p<0.05 is indicated by a single, p<0.01 double, and p<0.001 triple asterisk). Following the visual epoch, the encoding of spatial information becomes strongly significantly different between spiking and LFP signals. C. As in (B), the area under the tuning curve for each epoch and signal modality was computed, but after shifting each population's decoding values such that the decoding performance was 1 for the target in the preferred direction (i.e., Target 1). This measure of tuning width is only significantly different (paired t-test) between spiking and LFP activity during the delay period. D. Grid of statistical differences (paired t-test) in tuning

width across pairs of epochs computed separately for each signal modality. For spiking activity, the tuning width is only significantly different between the visual and motor epochs and between the late delay and motor epochs. For LFPs, the tuning width is significantly different across all epochs.

4.5 Discussion

In this study, we investigated the spatial tuning properties of spiking activity and local field potential signals in the SC, an oculomotor structure critical for the transformation of sensory input into motor commands. The combination of the anatomical organization of the SC and the typical electrophysiological approach lends itself to recording neural activity within a narrow column along the dorsoventral axis. Neurons within a column have largely similar preferred saccade directions as well as largely similar preferred visual target eccentricities (Gandhi & Katnani, 2011). We showed that despite this homogeneity, the active populations encode a wide range of directions. This population-level viewpoint provides insights into the spatial extent of direction tuning in the SC within a column that through single unit studies was thought to be essentially nonexistent for all visual angles except those close to the preferred direction.

For each short sliding window along the timeline of a delayed saccade task, a simple linear classifier was trained offline to categorize either spiking or LFP activity as belonging to one of eight target directions. We employed a classifier so that the amount of change in classification performance above baseline provided a singular measure of spatial information across the 24 channels on which activity was recorded. Such offline decoding algorithms have been used to characterize the spatial encoding properties of spiking activity (Boulay et al., 2016; Khanna et al., 2020; Ohmae et al., 2015) and LFP signals (Tremblay et al., 2015) in cortical oculomotor areas.

The implementation of classifiers to link neural activity to a behavioral phenomenon is beneficial because they provide a quantitative, comprehensive measurement of information encoding in neural populations (Glaser et al., 2020). Of note, we do not claim that the encoded information at any time represents a particular feature such as sensation, motor preparation, or motor initiation. Instead, we simply characterize the *amount* of information about direction present in the population throughout the timeline of sensorimotor integration. The end position of the saccade had to be within two degrees of the target position to count as a correct trial, which is a negligible displacement compared to the 45-degree angular distance between each pair of the eight targets used as the categories for classification. Thus, we have referred to the encoded target direction and saccade direction synonymously. However, a fine-scale characterization of the time points at which SC neurons encode spatial parameters in target-centered and gaze-centered coordinates can be found in Sajad et al., 2020.

Previous studies have compared the visual receptive fields of oculomotor neurons to their movement fields (equivalently, their spatial tuning properties during the respective visual and movement epochs). In cortical areas such as the FEF, the preferred target direction of individual neurons tends to be consistent between the visual and motor epochs (Khanna et al., 2020). The visual receptive fields of SC neurons have also been shown to largely overlap with their movement fields (Anderson et al., 1998; Wurtz & Goldberg, 1972). However, a more recent study demonstrated that despite this alignment, the movement fields are significantly larger than the visual receptive fields (Marino et al., 2008). Our results conform with these previous findings. When comparing the visual and motor epochs within a signal modality, we observed that the tuning width of spiking activity is significantly broader in the motor epoch than in the visual epoch (see Figure 28D).

Of much recent interest in the neuroscience community are the questions of what and how much information about various behavioral phenomena is contained in LFP signals – questions that have elicited studies on reach kinematic encoding by LFPs in primary motor cortex (Perel et al., 2015), attention in visual cortex (Prakash et al., 2021), route selection in hippocampus (Cheng et al., 2021), and grasping postures in anterior intraparietal cortex (Lehmann & Scherberger, 2015), among others. When comparing spatial encoding properties across the two simultaneously recorded signal modalities in this study, we found that the amount of spatial information present in spiking activity and LFPs diverged beginning in the early delay period, with the spike-based classifier better at decoding target location from small windows of activity. Both signal modalities displayed similar tuning widths to each other during the visual and motor epochs; in other words, the spatial extent of decoding performance across the eight targets was comparable between spike- and LFP-based classifiers during the visual response period and motor initiation period (considered independently – see Figure 28D). During the intervening delay period, the spike-based classifier performance remained high, but LFP-based performance dropped to near baseline levels for all targets except the target closest to the neural population's preferred direction. Thus, the encoding of direction is dynamic across epochs and signal modalities in the SC.

We suggest that the SC is a suitable candidate for brain-computer interface (BCI) applications, especially in BCIs implemented to address fundamental neuroscience questions (e.g., Sadtler et al., 2014). Although the vast majority of prior work that implements closed loop control of a computer cursor or robot arm has decoded neural activity from skeletomotor structures, a few groups have ventured into the oculomotor domain and demonstrated that volitional control of neural activity is possible in these areas (Graf & Andersen, 2014; Jia et al., 2017; Schafer & Moore, 2011) as well as in wholly non-motor areas (e.g., primary visual cortex, Neely et al., 2018). We foresee two possible limitations to using SC neurons or LFPs to decode intended saccade direction. First, the SC is a deep brain structure, which imposes a constraint on the number of recordable electrode sites. Cortical arrays fit electrode sites on the scale of hundreds, while laminar probes suitable for deep brain recording only allow for contacts on the order of tens. This is the likely reason that prior implementations of oculomotor BCIs have targeted cortical regions such as the lateral intraparietal area (LIP), frontal eye fields (FEF), and supplementary eye fields (SEF). However, advances in technology (e.g., Neuropixels) may soon negate this limitation. Second, the organization of neurons within a column along the dorsoventral axis results in neural populations with largely the same tuning properties (Gandhi & Katnani, 2011). This homogeneity theoretically reduces the spatial extent of decoding capability to targets far from the preferred target location, although we surprisingly observed that this is not the case; in fact, even targets in the diametrically opposite location of the preferred direction have above-chance decoding performance during the putatively preparatory delay period when the classifier is based on spiking activity (e.g., Figure 23). Nonetheless, a neural population with more varied preferred directions would maximize the spatial extent of high decoding performance. Recording from the FEF, a cortical oculomotor area, yields much more heterogeneity in directional tuning across electrode depth (Bruce et al., 1985), although due to its position in the bank of the arcuate sulcus the first limitation would still apply. Therefore, we are eager for the field to recognize the potential the SC has for brain-computer interface applications.

5.0 Summary and Conclusions

In this dissertation, I described three studies that add foundational knowledge of the role of the superior colliculus (SC) during active exploration of our visual environment. Each study provides a unique contribution to the characterization of SC population activity dynamics throughout various epochs of behavioral tasks that rely on sensorimotor integration – starting with the time course of visual- and motor-related signals during the period between sensation and action, continuing with the effect of behavioral context on the population activity patterns during the periods of sensation and action themselves, and finally ending with how SC populations encode broad ranges of visual stimulus angles and saccade directions throughout the entire timeline of sensorimotor integration through multiple signal modalities.

5.1 Summary of Chapter 2 – Timeline of Sensorimotor Transformation Signatures in the SC

In Chapter 2 we set out to characterize the dynamics of sensory and motor representations in the SC during the time period between the appearance of a visual stimulus and a goal-directed eye movement toward it. Although the SC is well known for its rapid firing rates signaling sensation and action on either end of sensorimotor integration, many neurons exhibit elevated activity during the intermediate period, which is often referred to as the "delay period" when a temporal delay is imposed on the sensation and action periods in a behavioral task. We employed a dimensionality reduction algorithm (refer to 1.3.1 for additional details about this machine learning technique) to view the correlated activity across each small SC population in a summarized fashion. First, we used a linear discriminant classifier (see 1.3.2 for an overview) to ensure that the population "subspaces" – that is, the patterns of activity represented in the low-dimensional framework (i.e., in Figure 2) – were separable between the sensory and action periods. Since we found that this was indeed the case, we computed the similarity of activity patterns during the delay period to either the visual or motor subspace using a probabilistic measure of the relative distance from either subspace. We observed a slow, systematic drift from a visual- to a motor-like representation as time in the delay period progressed. When small fixational saccades called microsaccades occurred at any point in the delay period, the population activity pattern transiently "reset" to a visual-like representation and was followed by a quick return to the original drifting dynamic. When we computed the correlation between the state-space position of the activity and the saccadic reaction time, a relationship emerged hundreds of milliseconds before the saccadic go cue.

These results prompt a number of future directions that could further build on this conceptualization of SC's role in sensorimotor transformation. First, an additional analysis of the "motor potential" of SC neurons at the end of the delay period would be fruitful. Previous work in the Gandhi lab has demonstrated that early disinhibition of the omnipause neurons (OPNs) in the brainstem often evokes a low-latency saccade (Jagadisan & Gandhi, 2016). This led to the conclusion that the SC does carry motor potential by the end of the delay period, and it resulted in another recently published study that provides evidence for "temporal stability" – in other words, the consistency in the population activity rank order across successive time points – as a potential mechanism for movement generation in the eye movement system. Building on the current study presented in Chapter 2, perhaps SC activity has higher motor potential (and thereby, a higher probability of inducing a saccade) on trials in which the population activity pattern has a higher motor-like representation as computed here. Both temporal stability and degree of motor-likeness may play a role in saccade initiation.

Historically, the large majority of oculomotor researchers have argued that the mechanisms for saccade generation are wholly distinct from those for arm reach generation due to the presence of OPNs acting as a gating signal, which are tonically active and are suppressed during a saccade (e.g., as noted in Darlington & Lisberger, 2020). It has also been thought that once SC populations reach a firing rate threshold, the OPNs are inhibited, and the SC motor command can reach the brainstem and the eye muscles for movement initiation (Jantz et al., 2013). Although this is an enticing view, it is overly simplistic. During the initial representation of a visual stimulus, the population activity often reaches the same firing rate level but does not evoke a movement. With this need for distinct representations of sensory and motor signals, it logically follows why we observed separable subspaces of population activity during the visual and motor epochs in the current study. Further studies are needed to solidify our understanding of the SC's role in movement initiation, but our findings suggest that movement generation mechanisms might be more conserved across motor systems than previously thought.

It would also be worthwhile to further investigate the effect of microsaccades on sensorimotor transformation. In this study, we grouped time periods of neural activity surrounding a detected microsaccade regardless of the time in the delay period at which the microsaccade occurred. These effects could be broken down by microsaccade timing (early or late) and direction (toward or away from the visual stimulus) to determine if the neural representation is unique across these conditions.

Finally, future studies could examine any change in the time course of sensorimotor transformation during tasks in which the cognitive load is higher. The memory-guided delayed saccade task (MG delay) is one such task – the sequence of a trial of this task is identical to the visually-guided delayed saccade task studied here with the exception of the disappearance of the visual stimulus shortly after its presentation. In the MG delay task, the animal must utilize working memory and make a delayed saccade to the remembered target location. Perhaps the population representation will evolve more rapidly or even "step" to a motor-like representation. On the other hand, the sensorimotor transformation signature may remain unaffected if cognitive factors such as this working memory component are encoded independently of sensory and motor features during the delay period (as observed in Chapter 3 during the periods of sensation and action). I hypothesize that microsaccades generated after target disappearance do not induce any transient perturbation in the population representation during the MG delay task because there is no longer peripheral visual stimulus to be jittered on the retina.

5.2 Summary of Chapter 3 – Contextual Effects on Sensorimotor Encoding in the SC

In the study presented in Chapter 3, we explored how neural populations in the SC encode two different forms of information – content (sensation/action) and context (delay/gap tasks) during the bookending visual and motor epochs of simple behavioral tasks. As in Chapter 2, we performed dimensionality reduction to summarize the pattern of activity across the population. The separability of the subspaces formed during the visual and motor epochs both within and across tasks were compared in a one-by-one manner using a simple linear discriminant classifier that was trained to identify activity patterns as belonging to one of two categories. The accuracy of the classifier above chance level gave an indication of the degree to which the activity patterns were distinct across conditions. The visual and motor patterns within each task were consistently separable, extending the result presented in Chapter 2 (i.e., that sensory and motor representations are distinct in the SC) to a second task. During both the sensory and motor periods, the population subspaces diverged across the two tasks, and we argued that the interaction of the bottom-up visual transient with movement preparation-related activity that begins accumulating in the gap period of the gap task (Edelman & Keller, 1996) serves to induce a different state in the oculomotor circuitry compared to trials of the delay task.

A potential future direction of this study is the exploration of population subspaces during a wider variety of behavioral and cognitive contexts. We started with a comparison of two tasks that have been thought to elicit similar neural responses in the SC but found that through a population-level analysis the patterns were surprisingly distinct across contexts. Therefore, it would be wise to examine the subspaces occupied by SC neural activity during a wider range of tasks, including the MG delay task described in 5.1. If the cognitive process of working memory adds complexity to the task, how are the population patterns affected? We have collected preliminary data in which neural populations were recorded while a monkey performed interleaved trials of the delay (i.e., visually guided, or VG), MG delay, and gap tasks, and performed initial analyses of the population representations during sensation and action. Having only a few data sets it is hard to form a concrete conclusion. I hypothesize that the sensory representation in the MG delay task will be overlapping that in the VG delay task because the trial flow is identical up to that point. Only after the stimulus disappears (i.e., after the transient sensory response is evoked) does the animal know that he is performing MG delay trial. I hypothesize that the motor representations will be slightly differentiable between the VG and MG delay tasks because of the different “neural state” imposed by the addition of a working memory component. We have already shown that if the activity is in a different initial state it will keep evolving differently even if the saccade produced has the same behavioral metrics (also see Afshar et al., 2011). We hope that future research explores the entire repertoire of SC network patterns and their multiplexed relationships to cognitive and behavioral phenomena.

5.3 Summary of Chapter 4 – Encoding of Spatial Parameters in Spiking and LFP SC Signals

In the final study, we explored the range of information about target/saccade direction contained within SC signals throughout the timeline of sensorimotor integration. Unlike in previous chapters where we only considered spiking activity, we included in this study an analysis of a second signal modality – the local field potential (LFP). We implemented a multiclass version of a linear discriminant analysis classifier (see overview in Offline Decoding 1.3.2), training it to categorize the target direction (one of eight possible) to which small windows of neural activity from single trials belonged. The performance at a particular time in the sensorimotor integration process gives an indication of the amount of information encoded by neural populations at that time. We found that the discriminability of target direction was rather broad and not limited to locations near the population’s preferred direction during the transient sensory and motor epochs for both signal modalities. For each signal modality, this encoding of direction by SC populations was also dynamic, exhibiting distinct tuning properties across the visual, delay, and motor epochs.

The logical next step for this line of research is investigating the time course of population-level information related to target/saccade *amplitude*, the other main spatial parameter encoded by

SC neurons. Here, we held the target amplitude constant across trials and varied the direction. Recording from an SC column while monkeys perform a delayed saccade task to a target with a constant direction across trials but with a varied amplitude would provide insights into this second dimension. To obtain a sufficient number of trials each recording session (equivalently, for each neural population), we limited the spatial resolution of both our behavioral paradigm and analyses to eight targets with 45-degree spacing increments. A variety of techniques are being developed in the field that seek to provide a way to combine neural activity across many sessions (Gallego et al., 2020; Pandarinath et al., 2018). We hope that in the future, the spatial resolution of spatial decoding can be implemented at a finer scale (i.e., with a narrower range of directions presented each day) and then combined across sessions through the use of such techniques.

In addition, I envision future analyses that break down the LFP signals into different frequency bands. For the analyses included in Chapter 4, all frequencies below 250 Hz were included in the time-domain LFP signal. The LFPs, at least the way quantified here, do not have much direction information outside of the visual and motor epochs. Do certain frequency bands yield better or worse classifiers? Do those frequency bands contain direction information that is obscured when all frequency bands are considered? A more thorough investigation of the relationship between population LFPs and target/saccade direction would be beneficial.

Finally, deep brain areas like the SC should be considered as a target structure for brain-computer interfaces (BCIs). Although deep brain structures are less accessible for electrophysiological recordings, the oculomotor system is arguably simpler than the skeletomotor system and is well-studied. The unique, structured organization of the SC in relation to spatial and sensorimotor parameters should make it an enticing structure of interest for future BCI experiments.

Bibliography

Afshar, A., Santhanam, G., Yu, B. M., Ryu, S. I., Sahani, M., & Shenoy, K. v. (2011). Single-trial neural correlates of arm movement preparation. *Neuron*, *71*(3), 555–564. https://doi.org/10.1016/j.neuron.2011.05.047

Anderson, R. W., Keller, E. L., Gandhi, N. J., & Das, S. (1998). Two-dimensional saccade-related population activity in superior colliculus in monkey. *Journal of Neurophysiology*, *80*(2), 798–817. https://doi.org/10.1152/JN.1998.80.2.798

Aoi, M. C., Mante, V., & Pillow, J. W. (2020). Prefrontal cortex exhibits multidimensional dynamic encoding during decision-making. *Nature Neuroscience*, *23*(11), 1410–1420. https://doi.org/10.1038/s41593-020-0696-5

Basso, M. A., & May, P. J. (2017). Circuits for Action and Cognition: A View from the Superior Colliculus. *Annual Review of Vision Science*, *3*, 197. https://doi.org/10.1146/ANNUREV-VISION-102016-061234

Bharmauria, V., Sajad, A., Yan, X., Wang, H., & Crawford, J. D. (2021). Spatiotemporal Coding in the Macaque Supplementary Eye Fields: Landmark Influence in the Target-to-Gaze Transformation. *ENeuro*, *8*(1), 1–23. https://doi.org/10.1523/ENEURO.0446-20.2020

Boulay, C. B., Pieper, F., Leavitt, M., Martinez-Trujillo, J., & Sachs, A. J. (2016). Single-trial decoding of intended eye movement goals from lateral prefrontal cortex neural ensembles. *Journal of Neurophysiology*, *115*, 486–499. https://doi.org/10.1152/jn.00788.2015.-Neu

Bruce, C. J., Goldberg, M. E., Bushnell, M. C., & Stanton, G. B. (1985). Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *Journal of Neurophysiology*, *54*(3), 714–734. https://doi.org/10.1152/JN.1985.54.3.714

Bryant, C. L., & Gandhi, N. J. (2005). Real-time data acquisition and control system for the measurement of motor and neural data. *Journal of Neuroscience Methods*, *142*(2), 193–200. https://doi.org/10.1016/J.JNEUMETH.2004.08.019

Buneo, C. A., Jarvis, M. R., Batista, A. P., & Andersen, R. A. (2002). Direct visuomotor transformations for reaching. *Nature*, *416*(6881), 632–636. https://doi.org/10.1038/416632A

Buonocore, A., Tian, X., Khademi, F., & Hafed, Z. M. (2021). Instantaneous movement-unrelated midbrain activity modifies ongoing eye movements. *ELife*, *10*. https://doi.org/10.7554/ELIFE.64150

Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, *13*(6), 407–420. https://doi.org/10.1038/nrn3241

Caruso, V. C., Mohl, J. T., Glynn, C., Lee, J., Willett, S. M., Zaman, A., Ebihara, A. F., Estrada, R., Freiwald, W. A., Tokdar, S. T., & Groh, J. M. (2018). Single neurons may encode simultaneous stimuli by switching between activity patterns. *Nature Communications*, *9*(1), 2715. https://doi.org/10.1038/s41467-018-05121-8

Caruso, V. C., Pages, D. S., Sommer, M. A., & Groh, J. M. (2018). Beyond the labeled line: Variation in visual reference frames from intraparietal cortex to frontal eye fields and the superior colliculus. *Journal of Neurophysiology*, *119*(4), 1411–1421. https://doi.org/10.1152/JN.00584.2017

Cheng, S., Li, M., Fan, J., Shang, Z., & Wan, H. (2021). Decoding route selection of pigeon during goal-directed behavior: A joint spike-LFP study. *Behavioural Brain Research*, *409*. https://doi.org/10.1016/J.BBR.2021.113289

Churchland, M. M., & Shenoy, K. v. (2007). Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *Journal of Neurophysiology*, *97*(6), 4235–4257. https://doi.org/10.1152/JN.00095.2007

Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., Bradley, D. C., Smith, M. A., Kohn, A., Movshon, A., Armstrong, K. M., Moore, T., Chang, S. W., Snyder, L. H., Lisberger, S. G., … Shenoy, K. v. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, *13*(3), 369–380. https://doi.org/10.1038/nn.2501

Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G., & Shenoy, K. v. (2006). Neural Variability in Premotor Cortex Provides a Signature of Motor Preparation. *Journal of Neuroscience*, *26*(14), 3697–3712. https://doi.org/10.1523/JNEUROSCI.3762-05.2006

Cooper, B., & McPeek, R. M. (2021). Role of the Superior Colliculus in Guiding Movements Not Made by the Eyes. *Annual Review of Vision Science*, *7*, 279–300. https://doi.org/10.1146/ANNUREV-VISION-012521-102314

Cowley, B. R., Kaufman, M. T., Butler, Z. S., Churchland, M. M., Ryu, S. I., Shenoy, K. v., & Yu, B. M. (2013). DataHigh: Graphical user interface for visualizing and interacting with high-dimensional neural activity. *Journal of Neural Engineering*, *10*(6), 066012. https://doi.org/10.1088/1741-2560/10/6/066012

Cowley, B. R., Smith, M. A., Kohn, A., & Yu, B. M. (2016). Stimulus-Driven Population Activity Patterns in Macaque Primary Visual Cortex. *PLoS*

Computational Biology, *12*(12), e1005185. https://doi.org/10.1371/JOURNAL.PCBI.1005185

Crapse, T. B., Lau, H., & Basso, M. A. (2018). A Role for the Superior Colliculus in Decision Criteria. *Neuron*, *97*(1), 181-194.e6. https://doi.org/10.1016/j.neuron.2017.12.006

Cunningham, J. P., & Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, *17*(11), 1500–1509. https://doi.org/10.1038/nn.3776

Darlington, T. R., & Lisberger, S. G. (2020). Mechanisms that allow cortical preparatory activity without inappropriate movement. *ELife*, *9*. https://doi.org/10.7554/eLife.50962

Das, V. E. (2016). Strabismus and the Oculomotor System: Insights from Macaque Models. *Annual Review of Vision Science*, *2*, 37–59. https://doi.org/10.1146/annurev-vision-111815-114335

Dekleva, B. M., Kording, K. P., & Miller, L. E. (2018). Single reach plans in dorsal premotor cortex during a two-target task. *Nature Communications*, *9*(1), 1–12. https://doi.org/10.1038/s41467-018-05959-y

Derya Übeyli, E. (2008). Analysis of EEG signals by combining eigenvector methods and multiclass support vector machines. *Computers in Biology and Medicine*, *38*(1), 14–22. https://doi.org/10.1016/J.COMPBIOMED.2007.06.002

Dorris, M. C., & Munoz, D. P. (1998). *Saccadic Probability Influences Motor Preparation Signals and Time to Saccadic Initiation*.

Duncker, L., & Sahani, M. (2021). Dynamics on the manifold: Identifying computational dynamical activity from neural population recordings. In *Current Opinion in Neurobiology* (Vol. 70, pp. 163–170). Elsevier Ltd. https://doi.org/10.1016/j.conb.2021.10.014

Dunn, C. A., Hall, N. J., & Colby, C. L. (2010). Spatial updating in monkey superior colliculus in the absence of the forebrain commissures: Dissociation between superficial and intermediate layers. *Journal of Neurophysiology*, *104*(3), 1267–1285. https://doi.org/10.1152/jn.00675.2009

Dyckman, K. A., Camchong, J., Clementz, B. A., & McDowell, J. E. (2007). An effect of context on saccade-related behavior and brain activity. *NeuroImage*, *36*(3), 774–784. https://doi.org/10.1016/J.NEUROIMAGE.2007.03.023

Edelman, J. A., & Goldberg, M. E. (2001). Dependence of Saccade-Related Activity in the Primate Superior Colliculus on Visual Target Presence. *Journal of Neurophysiology* 86(2), 676-691. www.jn.org

Edelman, J. A., & Keller, E. L. (1996). Activity of visuomotor burst neurons in the superior colliculus accompanying express saccades. *Journal of Neurophysiology*, *76*(2), 908–926. https://doi.org/10.1152/JN.1996.76.2.908

Everling, S., Dorris, M. C., Klein, R. M., & Munoz, D. P. (1999). *Role of Primate Superior Colliculus in Preparation and Execution of Anti-Saccades and Pro-Saccades*.

Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A., & Miller, L. E. (2020). Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience*, *23*(2), 260. https://doi.org/10.1038/S41593-019-0555-4

Gallego, J.A., Perich, M.G., Miller, L.E., Solla, S.A., 2017. Neural Manifolds for the Control of Movement. *Neuron* 94, 978–984. https://doi.org/10.1016/J.NEURON.2017.05.025

Gandhi, N. J., & Bonadonna, D. K. (2005). Temporal interactions of air-puff-evoked blinks and saccadic eye movements: insights into motor preparation. *Journal of Neurophysiology*, *93*(3), 1718–1729. https://doi.org/10.1152/JN.00854.2004

Gandhi, N. J., & Katnani, H. A. (2011). Motor functions of the superior colliculus. *Annual Review of Neuroscience*, *34*, 205–231. https://doi.org/10.1146/annurev-neuro-061010-113728

Gandhi, N. J., & Keller, E. L. (1999). Comparison of saccades perturbed by stimulation of the rostral superior colliculus, the caudal superior colliculus, and the omnipause neuron region. *Journal of Neurophysiology*, *82*(6), 3236–3253. https://doi.org/10.1152/JN.1999.82.6.3236

Glaser, J. I., Benjamin, A. S., Chowdhury, R. H., Perich, M. G., Miller, L. E., & Kording, K. P. (2020). Machine Learning for Neural Decoding. *ENeuro*, *7*(4), 1–16. https://doi.org/10.1523/ENEURO.0506-19.2020

Golub, M. D., Yu, B. M., & Chase, S. M. (2015). Internal models for interpreting neural population activity during sensorimotor control. *ELife, 4*. https://doi.org/10.7554/eLife.10015.001

Graf, A. B. A., & Andersen, R. A. (2014). Brain-machine interface for eye movements. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(49), 17630–17635. https://doi.org/10.1073/PNAS.1419977111

Hafed, Z. M. (2021). Superior colliculus saccade motor bursts do not dictate movement kinematics. *BioRxiv*, 2021.06.24.449726. https://doi.org/10.1101/2021.06.24.449726

Hafed, Z.M., Chen, C.Y., Tian, X., 2015. Vision, perception, and attention through the lens of microsaccades: Mechanisms and implications. *Frontiers in Systems Neuroscience* 9, 167. https://doi.org/10.3389/FNSYS.2015.00167/BIBTEX

Hafed, Z. M., & Krauzlis, R. J. (2010). Microsaccadic suppression of visual bursts in the primate superior colliculus. *Journal of Neuroscience*, *30*(28), 9542–9547. https://doi.org/10.1523/JNEUROSCI.1137-10.2010

Ikeda, T., Boehnke, S. E., Marino, R. A., White, B. J., Wang, C. A., Levy, R., & Munoz, D. P. (2015). Spatio-temporal response properties of local field potentials in the primate superior colliculus. *European Journal of Neuroscience, 41*(6), 856–865. https://doi.org/10.1111/EJN.12842

Ikeda, T., & Hikosaka, O. (2007). Positive and negative modulation of motor response in primate superior colliculus by reward expectation. *Journal of Neurophysiology*, *98*(6), 3163–3170. https://doi.org/10.1152/jn.00975.2007

Jagadisan, U. K., & Gandhi, N. J. (2016). Disruption of Fixation Reveals Latent Sensorimotor Processes in the Superior Colliculus. *Journal of Neuroscience*, *36*(22), 6129–6140. https://doi.org/10.1523/JNEUROSCI.3685-15.2016

Jagadisan, U. K., & Gandhi, N. J. (2017). Removal of inhibition uncovers latent movement potential during preparation. *ELife*, *6*. https://doi.org/10.7554/ELIFE.29648

Jagadisan, U. K., & Gandhi, N. J. (2022). Population temporal structure supplements the rate code during sensorimotor transformations. *Current Biology*, *32*. https://doi.org/10.1016/j.cub.2022.01.015

Jantz, J. J., Watanabe, M., Everling, S., & Munoz, D. P. (2013). Threshold mechanism for saccade initiation in frontal eye field and superior colliculus. *Journal of Neurophysiology*, *109*(11), 2767–2780. https://doi.org/10.1152/JN.00611.2012

Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, *13*(8), 1020. https://doi.org/10.1038/NN.2590

Jia, N., Brincat, S. L., Salazar-Gómez, A. F., Panko, M., Guenther, F. H., & Miller, E. K. (2017). Decoding of intended saccade direction in an oculomotor brain-computer interface. *Journal of Neural Engineering*, *14*(4), 046007. https://doi.org/10.1088/1741-2552/AA5A3E

Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. v. (2014). Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, *17*, 440–451. https://doi.org/10.1038/nn.3643

Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. v. (2015). Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *ELife*, *4*. https://doi.org/10.7554/ELIFE.04677

Keller, E. L., Lee, K. M., & McPeek, R. M. (2005). Readout of higher-level processing in the discharge of superior colliculus neurons. *Annals of the New York Academy of Sciences*, *1039*, 198–208. https://doi.org/10.1196/annals.1325.019

Khademi, F., Chen, C. Y., & Hafed, Z. M. (2020). Visual feature tuning of superior colliculus neural reafferent responses after fixational microsaccades. *Journal of Neurophysiology*, *123*(6), 2136–2153. https://doi.org/10.1152/JN.00077.2020/

Khanna, S. B., Scott, J. A., & Smith, M. A. (2020). Dynamic shifts of visual and saccadic signals in prefrontal cortical regions 8Ar and FEF. *Journal of Neurophysiology*, *124*(6), 1774–1791. https://doi.org/10.1152/JN.00669.2019

Kim, K., & Lee, C. (2017). Activity of primate V1 neurons during the gap saccade task. *Journal of Neurophysiology*, *118*(2), 1361–1375. https://doi.org/10.1152/JN.00758.2016

Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F., Qi, X. L., Romo, R., Uchida, N., & Machens, C. K. (2016). Demixed principal component analysis of neural population data. *ELife*, *5*. https://doi.org/10.7554/ELIFE.10989

Krauzlis, R. J., Lovejoy, L. P., & Zénon, A. (2013). Superior colliculus and visual spatial attention. In *Annual Review of Neuroscience* (Vol. 36, pp. 165–182). https://doi.org/10.1146/annurev-neuro-062012-170249

Lara, A. H., Elsayed, G. F., Zimnik, A. J., Cunningham, J. P., & Churchland, M. M. (2018). Conservation of preparatory neural events in monkey motor cortex regardless of how movement is initiated. *ELife*, *7*. https://doi.org/10.7554/eLife.31826.001

Latimer, K. W., Yates, J. L., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, *349*(6244), 184–187. https://doi.org/10.1126/science.aaa4056

Lee, J., & Groh, J. M. (2012). Auditory signals evolve from hybrid- to eye-centered coordinates in the primate superior colliculus. *Journal of Neurophysiology*, *108*(1), 227–242. https://doi.org/10.1152/JN.00706.2011

Lehmann, S. J., & Scherberger, H. (2015). Spatial Representations in Local Field Potential Activity of Primate Anterior Intraparietal Cortex (AIP). *PloS One*, *10*(11). https://doi.org/10.1371/JOURNAL.PONE.0142679

Leon, M. I., & Shadlen, M. N. (2003). Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron*, *38*(2), 317–327. https://doi.org/10.1016/S0896-6273(03)00185-5

Libby, A., & Buschman, T. J. (2021). Rotational dynamics reduce interference between sensory and memory representations. *Nature Neuroscience*, *24*, 715–726. https://doi.org/10.1038/s41593-021-00821-9

Lovejoy, L. P., & Krauzlis, R. J. (2010). Inactivation of primate superior colliculus impairs covert selection of signals for perceptual judgments. *Nature Neuroscience*, *13*(2), 261–266. https://doi.org/10.1038/nn.2470

Marino, R. A., Levy, R., & Munoz, D. P. (2022). Distinct sensory- and goal-related signals underlie the gap effect in the superior colliculus. *European Journal of Neuroscience*, *55*(1), 205–226. https://doi.org/10.1111/ejn.15533

Marino, R. A., Rodgers, C. K., Levy, R., & Munoz, D. P. (2008). Spatial relationships of visuomotor transformations in the superior colliculus map. *Journal of Neurophysiology*, *100*(5), 2564–2576. https://doi.org/10.1152/JN.90688.2008

Martinez-Conde, S., Macknik, S. L., & Hubel, D. H. (2004). The role of fixational eye movements in visual perception. *Nature Reviews Neuroscience*, *5*(3), 229–240. https://doi.org/10.1038/nrn1348

Massot, C., Jagadisan, U. K., & Gandhi, N. J. (2019). Sensorimotor transformation elicits systematic patterns of activity along the dorsoventral extent of the superior colliculus in the macaque monkey. *Communications Biology*, *2*(1), 1–14. https://doi.org/10.1038/s42003-019-0527-y

McPeek, R. M., & Keller, E. L. (2002). Saccade Target Selection in the Superior Colliculus During a Visual Search Task. *Journal of Neurophysiology*, 88(4), 2019-2034. https://doi.org/10.1152/jn.00181.2002

Miyashita, N., & Hikosaka, O. (1996). Minimal synaptic delay in the saccadic output pathway of the superior colliculus studied in awake monkey. *Experimental Brain Research*, *112*(2), 187–196. https://doi.org/10.1007/BF00227637

Mohler, C. W., & Wurtz, R. H. (1976). Organization of monkey superior colliculus: intermediate layer cells discharging before eye movements. *Journal of Neurophysiology*, *39*(4), 722–744. https://doi.org/10.1152/JN.1976.39.4.722

Mullette-Gillman, O. A., Cohen, Y. E., & Groh, J. M. (2005). Eye-centered, head-centered, and complex coding of visual and auditory targets in the intraparietal sulcus. *Journal of Neurophysiology*, *94*(4), 2331–2352. https://doi.org/10.1152/JN.00021.2005

Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M., & Carmena, J. M. (2018). Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. *Neuron*, *97*(6), 1356-1368.e4. https://doi.org/10.1016/J.NEURON.2018.01.051

Ohmae, S., Takahashi, T., Lu, X., Nishimori, Y., Kodaka, Y., Takashima, I., & Kitazawa, S. (2015). Decoding the timing and target locations of saccadic eye movements from neuronal activity in macaque oculomotor areas. *Journal of Neural Engineering*, *12*(3), 036014. https://doi.org/10.1088/1741-2560/12/3/036014

Pandarinath, C., O'Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., Henderson, J. M., Shenoy, K. v., Abbott, L. F., & Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods 2018 15:10*, *15*(10), 805–815. https://doi.org/10.1038/s41592-018-0109-9

Peel, T. R., Dash, S., Lomber, S. G., & Corneil, B. D. (2021). Frontal eye field inactivation alters the readout of superior colliculus activity for saccade generation in a task-dependent manner. *Journal of Computational Neuroscience*, *49*(3), 229–249. https://doi.org/10.1007/S10827-020-00760-7

Peel, T. R., Hafed, Z. M., Dash, S., Lomber, S. G., & Corneil, B. D. (2016). A Causal Role for the Cortical Frontal Eye Fields in Microsaccade Deployment. *PLoS Biology*, *14*(8). https://doi.org/10.1371/JOURNAL.PBIO.1002531

Perel, S., Sadtler, P. T., Oby, E. R., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., & Chase, S. M. (2015). Single-unit activity, threshold crossings, and local field potentials in motor cortex differentially encode reach kinematics. *Journal of Neurophysiology*, *114*(3), 1500. https://doi.org/10.1152/JN.00293.2014

Prakash, S. S., Das, A., Kanth, S. T., Mayo, J. P., & Ray, S. (2021). Decoding of Attentional State Using High-Frequency Local Field Potential Is As Accurate As Using Spikes. *Cerebral Cortex*, *31*(9), 4314–4328. https://doi.org/10.1093/CERCOR/BHAB088

Robinson, D. L., & Wurtz, R. H. (1976). Use of an extraretinal signal by monkey superior colliculus neurons to distinguish real from self-induced stimulus movement. *Journal of Neurophysiology*, *39*(4), 852–870. https://doi.org/10.1152/JN.1976.39.4.852

Sadeh, M., Sajad, A., Wang, H., Yan, X., Crawford, J.D., 2020. Timing Determines Tuning: A Rapid Spatial Transformation in Superior Colliculus Neurons during Reactive Gaze Shifts. *eNeuro* 7. https://doi.org/10.1523/ENEURO.0359-18.2019

Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., & Batista, A. P. (2014). Neural constraints on learning. *Nature*, *512*(7515), 423. https://doi.org/10.1038/NATURE13665

Sajad, A., Sadeh, M., & Crawford, J. D. (2020). Spatiotemporal transformations for gaze control. *Physiological Reports*, *8*(16), e14533. https://doi.org/10.14814/PHY2.14533

Sajad, A., Sadeh, M., Yan, X., Wang, H., & Crawford, J. D. (2016). Transition from Target to Gaze Coding in Primate Frontal Eye Field during Memory Delay and Memory–Motor Transformation. *ENeuro*, *3*(2), 82. https://doi.org/10.1523/ENEURO.0040-16.2016

Schafer, R. J., & Moore, T. (2011). Selective Attention from Voluntary Control of Neurons in Prefrontal Cortex. *Science*, *332*(6037), 1568. https://doi.org/10.1126/SCIENCE.1199892

Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. *Neuron*, *102*(1), 249-259.e4. https://doi.org/10.1016/J.NEURON.2019.01.026

Shen, K., Valero, J., Day, G. S., & Paré, M. (2011). Investigating the role of the superior colliculus in active vision with the visual search paradigm. *European Journal of Neuroscience*, *33*(11), 2003–2016. https://doi.org/10.1111/j.1460-9568.2011.07722.x

Shenoy, K. v, Sahani, M., & Churchland, M. M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, *36*, 337–359. https://doi.org/10.1146/annurev-neuro-062111-150509

Smalianchuk, I., Jagadisan, U. K., & Gandhi, N. J. (2018). Instantaneous Midbrain Control of Saccade Velocity. *The Journal of Neuroscience*, *38*(47), 10156. https://doi.org/10.1523/JNEUROSCI.0962-18.2018

Sommer, M. A., & Wurtz, R. H. (2000). Composition and topographic organization of signals sent from the frontal eye field to the superior colliculus. *Journal of Neurophysiology*, *83*(4), 1979–2001. https://doi.org/10.1152/JN.2000.83.4.1979/

Sparks, D. L. (1999). Conceptual issues related to the role of the superior colliculus in the control of gaze. *Current Opinion in Neurobiology*, *9*(6), 698–707.

Sparks, D., Rohrer, W. H., & Zhang, Y. (2000). The role of the superior colliculus in saccade initiation: a study of express saccades and the gap effect. In *Vision Research* (Vol. 40). www.elsevier.com

Tanaka, M. (2007). Cognitive Signals in the Primate Motor Thalamus Predict Saccade Timing. *Journal of Neuroscience*, *27*(44), 12109–12118. https://doi.org/10.1523/JNEUROSCI.1873-07.2007

Trautmann, E. M., Stavisky, S. D., Lahiri, S., Ames, K. C., Kaufman, M. T., O'Shea, D. J., Vyas, S., Sun, X., Ryu, S. I., Ganguli, S., & Shenoy, K. v. (2019). Accurate Estimation of Neural Population Dynamics without Spike Sorting. *Neuron*, *103*(2), 292-308.e4. https://doi.org/10.1016/J.NEURON.2019.05.003

Tremblay, S., Doucet, G., Pieper, F., Sachs, A., & Martinez-Trujillo, J. (2015). Single-Trial Decoding of Visual Attention from Local Field Potentials in the Primate Lateral Prefrontal Cortex Is Frequency-Dependent. *Journal of Neuroscience*, *35*(24), 9038–9049. https://doi.org/10.1523/JNEUROSCI.1041-15.2015

Wurtz, R. H., & Goldberg, M. E. (1972). Activity of superior colliculus in behaving monkey. 3. Cells discharging before eye movements. *Journal of Neurophysiology*, *35*(4), 575–586. https://doi.org/10.1152/JN.1972.35.4.575

Wurtz, R. H., & Optican, L. M. (1994). Superior colliculus cell types and models of saccade generation. *Current Opinion in Neurobiology*, *4*(6), 857–861. https://doi.org/10.1016/0959-4388(94)90134-1

Wurtz, R. H., Sommer, M. A., Paré, M., & Ferraina, S. (2001). Signal transformations from cerebral cortex to superior colliculus for the generation of saccades. *Vision Research*, *41*(25–26), 3399–3412. https://doi.org/10.1016/S0042-6989(01)00066-9

Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. v., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*(1), 614–635. https://doi.org/10.1152/JN.90941.2008/

Zhi, Y. X., Lukasik, M., Li, M., Dolatabadi, E., Wang, R. H., & Taati, B. (2018). Automatic Detection of Compensation During Robotic Stroke Rehabilitation Therapy. *IEEE Journal of Translational Engineering in Health and Medicine*, *6*. https://doi.org/10.1109/JTEHM.2017.2780836