THE UNIVERSITY OF CHICAGO

CORTICAL AND SUBCORTICAL CONTRIBUTIONS TO VISUAL CATEGORY DECISIONS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES
AND THE PRITZKER SCHOOL OF MEDICINE
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL NEUROSCIENCE

BY
BARBARA PEYSAKHOVICH

CHICAGO, ILLINOIS
JUNE 2023

For [Bab, Mam, Ded]ulja

"The eye obeys exactly the action of the mind. When a thought strikes us, the eyes fix, and remain gazing at a distance; in enumerating the names of persons or of countries, as France, Germany, Spain, Turkey, the eyes wink at each new name. There is no nicety of learning sought by the mind, which the eyes do not vie in acquiring."

- Ralph Waldo Emerson, *The Conduct of Life*

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I am tremendously grateful to my advisor, Dave Freedman, for his guidance and mentorship over the past six years. Dave has been unwavering in his support and encouragement, and I have learned so much from Dave's careful and creative approach to science. I would also like to thank past and current members of the Freedman Lab. It's been a true joy to work with such bright, kind, and fun people, and I will miss our science chats and lunchtime BeaverGang tournaments.

I am deeply grateful to my thesis committee, Jason MacLean, John Maunsell, and Matt Kaufman, for their insightful comments and valuable feedback on my research throughout the course of my PhD. Beyond helping to shape my PhD project, each member of my committee has provided a model for how to be a thoughtful, rigorous scientist, and they will continue to be a source of inspiration for the rest of my career.

I would also like to thank my incredible family. Babulja, Mamulja, Dedulja, Rob, Richard, Marie and David: Thank you for filling my life with love. I feel your unconditional support every day, even when we are oceans apart. Thank you to Zoya for being the brightest light in my life, and to Bernie for being a reliable source of love, warmth, and tail wags. Finally, I would like to thank Tony, who has been my rock for the past 10+ years. Thank you for moving across oceans with me as I pursued my dreams, for being a constant source of encouragement, and for being the best parent to Zoya (and Bernie). I am forever grateful for your love and support.

ABSTRACT

Categorization is a fundamental cognitive process by which the brain assigns stimuli to behaviorally meaningful groups. Previous studies on visual categorization in primates have identified a hierarchy of cortical areas that are involved in the transformation of veridical sensory information into abstract category representations. However, categorization behaviors are ubiquitous across diverse animal species, even those without a neocortex, motivating the possibility that subcortical regions may contribute to abstract cognition in primates. One candidate structure is the superior colliculus (SC), a midbrain region that is evolutionarily conserved across vertebrates. Although traditionally thought to mediate only reflexive spatial orienting behaviors, especially saccades in primates, the SC is also involved in cognitive tasks that require spatial orienting.

In the first part of this thesis, we investigated the involvement of the primate SC in abstract categorization, and show that the SC plays an unexpected key role in higher-order, non-spatial cognition. We trained monkeys to group motion stimuli into categories based on an arbitrary rule, and compared neural activity in the SC and the lateral intraparietal area (LIP), a cortical region previously shown to causally contribute to category decisions, while monkeys performed this task. We observed unexpectedly strong and short-latency category encoding in the SC that was more reliable and arose even earlier than in the LIP. Moreover, monkeys' performance on the categorization task was markedly impaired during reversible inactivation of the SC, indicating that the observed category signals in the SC may causally contribute to category processing. In addition, we show that category and eye movement-related signals are encoded in near-orthogonal subspaces in population activity in the SC, providing an explanation for how a motor structure like the SC can be recruited to participate in more flexible cognitive behaviors. These results extend the well-established role of the SC in spatial orienting to non-spatial, higher-order cognition.

In the second part of this thesis, we investigated how behavioral task demands affect

category and sensory encoding in the SC, LIP, and the middle temporal area, (MT), an early visual cortical area that is involved in motion processing. We trained monkeys to alternate between blocks of the motion categorization task and blocks in which they passively viewed the same stimuli and received a reward for maintaining fixation. The physical stimulus and stimulus location was identical in the two blocks, but only the categorization task required the monkeys to use the stimulus information to obtain a reward; therefore, we could compare, in the same neurons, how behavioral context affects stimulus encoding. We observed significantly weaker stimulus direction encoding during passive viewing than during the categorization task in all three brain areas. Moreover, although both LIP and SC encoded stimulus category during the categorization task, category encoding was largely absent during passive viewing in both areas. These results indicate that neural populations in LIP and SC can flexibly route sensory input based on current behavioral demands.

CHAPTER 1

INTRODUCTION

Categorization is a fundamental cognitive process that enables animals to efficiently process and organize information in their environment, thus facilitating their understanding and memory of the world and allowing them to form decisions, make predictions, and adapt to novel situations. Animals (especially primates) are remarkably skilled at rapidly categorizing complex stimuli into behaviorally relevant groups. How do our brains perform this (seemingly effortless) feat of flexibly transforming low-level sensory information into high-level representations of categories?

In advanced animals, category rules are often learned from experience, and categories can be defined by abstract features rather than shared physical characteristics. Because the higher-order cognitive processing required for this kind of abstract categorization has traditionally been attributed to the neocortex, previous investigations of the neural mechanisms of abstract categorization have largely focused on cortical regions. These studies have identified several cortical brain regions in the rhesus macaque that encode categories of visual stimuli during abstract categorization tasks, including the prefrontal cortex (Freedman et al., 2001; Swaminathan and Freedman, 2012; Antzoulatos and Miller, 2011), the lateral intraparietal area (Freedman and Assad, 2006; Sarma et al., 2016; Rishel et al., 2013; Mohan et al., 2021), and the frontal eye fields (Ferrera et al., 2009).

In an early study, monkeys were trained to categorize parametrically morphed "cat" and "dog" stimuli. Activity of neurons in the prefrontal cortex (PFC) reflected the category membership of the stimuli; many PFC neurons preferentially responded to all stimuli in one of the two categories, even though the category was comprised of stimuli that varied considerably in their physical features (Freedman et al., 2001). In more recent studies, monkeys were trained to categorize motion stimuli that consist of videos of a patch of dots moving coherently in a particular direction. In these experiments, 360° of motion directions

were assigned to two categories based on an arbitrary category boundary (**Fig. 1.1a**). After animals learn the motion categorization task, neurons in LIP acquire category selectivity (i.e., preferentially respond to all motion stimuli that belong to one category over all stimulus in the other category; **Fig. 1.1b**). (Freedman and Assad, 2006). However, experience with the categorization task does not affect affect responses across all visual cortical areas. Neurons in the middle temporal area (MT), an extrastriate region that receives direct input from the primary visual cortex (V1) (Cragg, 1969) and is involved processing of motion stimuli (Maunsell and Van Essen, 1983b), remained tuned to direction rather than category even after extensive training on the category task (Freedman and Assad, 2006).

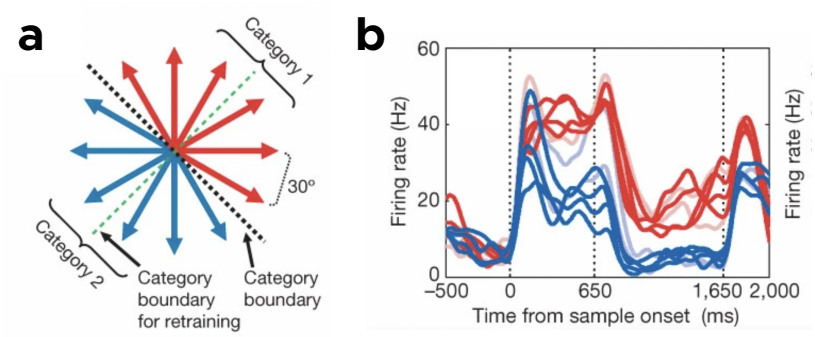

Figure 1.1: Neurons in the lateral intraparietal area signal learned categories. a, Schematic illustrating how 360° of motion directions are assigned to two groups based on an arbitrary category rule. **b,** An example category-tuned neuron from the lateral intraparietal area (LIP). The neuron preferentially responds to motion directions that belong to one of the two categories. Adapted from (Freedman and Assad, 2006)

Converging evidence suggests that the LIP may play a central role in visual categorization. In a study that compared neuronal responses in the LIP and PFC during motion categorization, LIP neurons had stronger, more predictive, and shorter-latency category signals than PFC neurons (Swaminathan and Freedman, 2012). Moreover, a recent experiment reported that reversible inactivation of the LIP caused significant impairments in performance on a motion categorization task when stimuli were placed in the inactivated hemifield (Zhou and Freedman, 2019), indicating that the LIP may be causally involved in sensory evaluation of stimuli during categorization.

In this thesis, we extended these investigations of the neural basis of visual motion categorization to the superior colliculus (SC), a midbrain structure that is involved spatial orienting and target selection. The LIP and the SC are both core nodes of the primate oculomotor network, and the two brain regions are reciprocally connected through direct descending projections from the LIP to the SC and indirect ascending projects from the SC to the LIP through the pulvinar nucleus in the thalamus (Fries, 1984; Lynch et al., 1985; Asanuma et al., 1985; Andersen et al., 1990; Paré and Wurtz, 1997; Clower et al., 2001).

1.1 The role of the superior colliculus in orienting movements

The superior colliculus (or tectum) is a brain structure located on the dorsal surface of the midbrain that evolutionarily conserved across all vertebrate species. The SC has been studied for more than a century in the context of its role in directing orienting movements of the eyes and head. The earliest experiments that linked the SC to eye movements showed that electrical stimulation of the SC generates reliable eye movements (Adamük, 1870). Decades later, electrophysiological recordings in rhesus macaques revealed that many neurons in the SC, particularly in the deep and intermediate layers, reliably discharge before (and during) eye movements made in a particular range of directions (Schiller and Koerner, 1971; Schiller and Stryker, 1972; Wurtz and Goldberg, 1971, 1972; Sparks, 1978). The role of the primate SC in generating eye movements was confirmed through causal perturbations that showed that (1) electrical stimulation of the SC in macaques generates movements of the eye and neck (Robinson, 1972; Schiller and Stryker, 1972; Van Opstal et al., 1990; Freedman et al., 1996; Stanford et al., 1996), and that (2) application of muscimol or lidocaine to reversibly inactivate the SC significantly reduces velocity and amplitude of eye movements contralateral to the inactivation hemifield (Hikosaka and Wurtz, 1985, 1986; Lee et al., 1988), decreases accuracy of contralateral saccades (especially for saccades to remembered locations) (Hikosaka and Wurtz, 1985), and increases latency of contralateral saccades (Hikosaka and

Wurtz, 1985, 1986; Schiller et al., 1987).

This large body of work, spanning more than a century, established the SC as a key brain region involved in coordinating and executing eye movements and established a clear causal link between activity in the SC and saccadic eye movements. For a long time, the SC was thought to be a low-level, reflexive structure that merely receives and implements instructions from upstream brain areas where to orient to.

1.2 The role of the superior colliculus in target selection

Some of the first experiments to indicate that the SC may be involved in the deliberation of what kind of saccade to make showed that activity of neurons in the intermediate layer of the SC are differentially modulated depending on the certainty of the impending saccade, when the parameters of the saccades are similar (Basso and Wurtz, 1997, 1998). Almost a decade later, (McPeek and Keller, 2004) provided one of the first causal reports that the SC is involved in saccade target selection, rather than just motor command implementation. In this study, monkeys performed a color-search task in which the they viewed an array of four visual targets that included three distractors that matched in color and one color-mismatched target. The monkeys were trained to identify and saccade to the odd-one-out target in order to receive a reward. The authors reversibly inactivated the SC and compared animals' performance on the task before and during SC inactivation. The key manipulation in the experiment was the inclusion of two conditions of varying difficulty: in the easier condition, there was a large perceptual difference between the color of the target and distractors, and in the difficult condition, the perceptual difference between the target and distractor was small. If the role of the SC is limited to simply generating a saccade based on instructions from another brain region, inactivation of the SC should produce similar impairments in performance on the easy and difficult task conditions. However, if the SC is causally involved in choosing which of the four targets to make a saccade to, then

inactivation of the SC should cause larger impairments on the difficulty condition than the easier condition. Indeed, SC inactivation caused significantly greater impairments on the difficult task condition, indicating that the SC is causally involved in the choice of where to saccade to.

1.3 The role of the SC in cognitive spatial processing

More recent work has shown that the SC is involved in even more complex types of spatial orienting functions (Basso et al., 2021). In particular, the SC has been shown to play a causal role in the deployment of covert spatial attention (Krauzlis et al., 2014, 2013), as well as decision making during tasks in which animals report their choice with a saccade.

1.3.1 The primate SC in involved in the deployment of spatial attention

Converging evidence indicate that the primate SC is a key brain area involved in directing spatial attention to locations of interest (Krauzlis et al., 2014, 2013). Spatial attention has been studied using attentional cuing tasks in which animals are directed to covertly orient (or attend) to a cued spatial location without moving their eyes. For example, in the classical change-detection task, monkeys receive a cue indicating the spatial location that they will need to attend to later in the trial. After the cue disappears, two stimuli appear, and the monkeys are trained to indicate when there is a change in the stimulus at the previously cued location (and ignore changes of the stimulus at the un-cued location).

Studies using this (or similar) paradigms have shown that neurons in the SC signal the location of covert attention; SC neurons can have vastly different responses to identical stimuli depending on whether the animal is attending to that stimulus or to a stimulus at another location (Ignashchenkova et al., 2004). Causal manipulations have provided further evidence that the SC is critically involved in the deployment of spatial attention. Low-intensity electrical stimulation of the SC enhances behavioral detection (Cavanaugh and

Wurtz, 2004; Cavanaugh et al., 2006) and discrimination of motion directions, selectively for stimuli presented in the spatial locations that correspond with the receptive fields of the stimulated region of SC (Müller et al., 2005; Cavanaugh and Wurtz, 2004). Moreover, inactivation of the SC impairs animals' ability to attend to stimuli in the affected spatial location (Zénon and Krauzlis, 2012), though only when there is a distractor stimulus present in an unaffected location (Lovejoy and Krauzlis, 2010).

1.3.2 The primate SC is involved in sensorimotor decision-making

The role of the SC in more cognitive spatial functions also extends to decision-making tasks in which animals report their choices with a saccade, such as the perceptual decision-making paradigm, a classical style of task that has been used to study decision-making processes in primates for decades (Britten et al., 1992; Gold and Shadlen, 2007). In these tasks, monkeys are trained to report their perceptual decisions about noisy or ambiguous stimuli. In the popular random-dot-motion variant of this task, monkeys view visual stimuli that consist of a patch of moving dots, some proportion of which (as determined by the experimenter) are moving coherently in a particular direction while the rest are moving randomly. The animals are tasked with making a decision about this (mostly) noisy stimulus (on average, are the dots moving more to the left or to the right?) and making a saccade to the target associated with their choice. The experimenters can control the difficulty of the decision on each trial by varying the coherence of the stimulus. The noisiness of the stimuli is a critical aspect of the task; the "decision" component is the weighing and accumulation of the sensory evidence that culminates in the decision of which target to orient to.

A large number of studies have shown that the SC is important for forming and representing sensorimotor decisions during. Neurons in the SC have been shown to encode animals' decisions these types of perceptual decision tasks, and their firing rates are modulated by the coherence of the stimuli (Horwitz and Newsome, 1999, 2001; Ratcliff et al., 2003; Hor-

witz et al., 2004; Ratcliff et al., 2007). In a task designed to systematically shift monkeys' decision criteria, activity of SC neurons has been shown to track these shifts in criterion (Crapse et al., 2018). Finally, recent studies studies have shown that inactivation of the SC disrupts the accumulation of evidence towards choices associated with targets in the affected locations (Jun et al., 2021; Stine et al., 2022), indicating that the SC is causally involved in sensorimotor decisions.

Together, these studies provide strong support for the notion that the SC is not just a reflexive orienting structure, but rather is involved in the deliberation of where to orient to. The role of the SC in target selection even extends to more complex tasks, such as the attentional cuing and perceptual decision paradigms described above. However, those studies used tasks in which animals reported their decisions with eye movements, or tasks that directed animals to covertly orient to stimuli at particular spatial locations. It is therefore unknown whether the involvement of the SC in more cognitive tasks is restricted to contexts involving orienting, either overt (through saccades) or covert (through modulation of spatial attention). In the next chapter of this thesis, we investigated whether the SC is more generally involved in non-spatial cognition during an abstract visual motion categorization task.

CHAPTER 2

PRIMATE SUPERIOR COLLICULUS IS ENGAGED IN ABSTRACT HIGHER-ORDER COGNITION

Barbara Peysakhovich, Stephanie M. Tetrick, Alessandra A. Silva, Sihai Li, Ou Zhu, Guilhem Ibos, W. Jeffrey Johnston, and David J. Freedman

Author contributions: BP, GI and DJF conceived of the project. BP supervised animal training, and BP, SMT, AAS, OZ, GI, and WJJ performed animal training and handling. BP and SMT collected electrophysiological recording data. BP collected inactivation data with assistance from AAS and SL. BP performed all data analysis and visualization. BP and DJF wrote the manuscript with input from AAS, SL, OZ, GI, and WJJ.

Abstract

The superior colliculus (SC) is an evolutionarily conserved midbrain region that is traditionally thought to mediate spatial orienting, including saccadic eye movements and covert spatial attention. Here, we reveal a novel role of the SC in abstract cognition, independent of its role in spatial orienting. We compared neural activity in the primate SC and the lateral intraparietal area (LIP), a cortical region previously shown to causally contribute to category decisions, during an abstract visual categorization task. We found that the SC exhibits stronger and shorter-latency encoding of learned categories than the LIP. Moreover, reversible pharmacological inactivation of the SC markedly caused marked impairments in animals' performance on the categorization task. These results demonstrate that the SC mediates abstract, higher-order cognitive processes that have traditionally been attributed to the neocortex.

Keywords: superior colliculus, cognition, categorization, subcortical, vision, parietal, decision making, primate, electrophysiology, reversible inactivation

2.1 Introduction

The superior colliculus (SC), a brainstem region that is evolutionarily conserved across vertebrates, has long been known to play a crucial role in directing orienting movements of the eyes and head (Adamük, 1870; Apter, 1946; Wurtz and Goldberg, 1971; Schiller and Koerner, 1971; Schiller and Stryker, 1972; Robinson, 1972; Wurtz and Goldberg, 1972; Sparks, 1978; Hikosaka and Wurtz, 1985, 1986; Lee et al., 1988). Although traditionally thought to be a reflexive structure that receives and implements motor instructions from upstream brain areas, the SC is involved in the selection of which target to orient to (Basso and Wurtz, 1997, 1998; McPeek and Keller, 2004; Carello and Krauzlis, 2004), and has been shown to be engaged in complex behavioral tasks that involve either overt target selection, such as decision-making tasks in which animals report their choice with a saccade to a particular target (Horwitz and Newsome, 1999, 2001; Ratcliff et al., 2003; Horwitz et al., 2004; Ratcliff et al., 2007; Crapse et al., 2018; Jun et al., 2021; Stine et al., 2022), or covert target selection, such as tasks that require animals to attend to stimuli at particular spatial locations (Ignashchenkova et al., 2004; Cavanaugh and Wurtz, 2004; Müller et al., 2005; Cavanaugh et al., 2006; Lovejoy and Krauzlis, 2010; Zénon and Krauzlis, 2012; Krauzlis et al., 2013, 2014). However, it is unknown whether the involvement of the SC in cognitively demanding tasks is restricted to contexts involving spatial orienting or target selection.

Here, we investigated whether the primate SC is more generally involved in abstract cognition. We trained monkeys to perform an abstract visual categorization task that dissociates sensory, cognitive, and motor components, and compared neuronal activity in the SC and the lateral intraparietal area (LIP), a cortical region in the posterior parietal cortex that is anatomically interconnected with the SC (Fries, 1984; Lynch et al., 1985; Asanuma et al., 1985; Andersen et al., 1990; Paré and Wurtz, 1997; Clower et al., 2001) and has been previously shown to causally contribute to category processing (Zhou and Freedman, 2019). We also reversibly inactivated the SC to assess its causal contribution to category decisions.

We show that the SC exhibits robust, short-latency encoding of abstract categories and that inactivation of the SC markedly impairs animals' categorization task performance. These results indicate that the primate SC plays an unexpected key role in higher-order cognition, independent of its role in spatial orienting. In addition, we show that category and saccade-related signals are encoded in near-orthogonal subspaces in population activity in the SC, providing an explanation for how a motor structure like the SC can be recruited to participate in more flexible cognitive behaviors.

2.2 Results

2.2.1 Behavior

Two monkeys performed a delayed match-to-category (DMC) task in which they grouped 360° of motion directions into two categories based on an arbitrary category rule. The categories were defined by two perpendicular boundaries that produced four 90°-wide quadrants (**Fig. 2.1a**). To disambiguate neuronal encoding of direction vs. category, opposite quadrants were assigned to the same category, so that motion directions that are 180° apart belonged to the same category while nearby directions were often in different categories. Monkeys viewed sample and test dot-motion stimuli separated by a 1.2 sec memory delay (**Fig. 2.1b**). On each trial, they received a fluid reward for releasing a manual touch bar when the category of the test stimulus matched the category of the sample stimulus (Match trials). If the test stimulus category did not match the sample stimulus category (Non-match trials), the monkeys were shown a second test stimulus, which always matched the sample category (and required release of the touch bar in order to receive a reward). The monkeys were required to maintain gaze on a central fixation spot throughout the trial (see *Methods*). The monkeys' decisions about the sample category were abstract because the two categories were defined by the learned arbitrary boundaries, and because they were not

linked to different motor actions or plans.

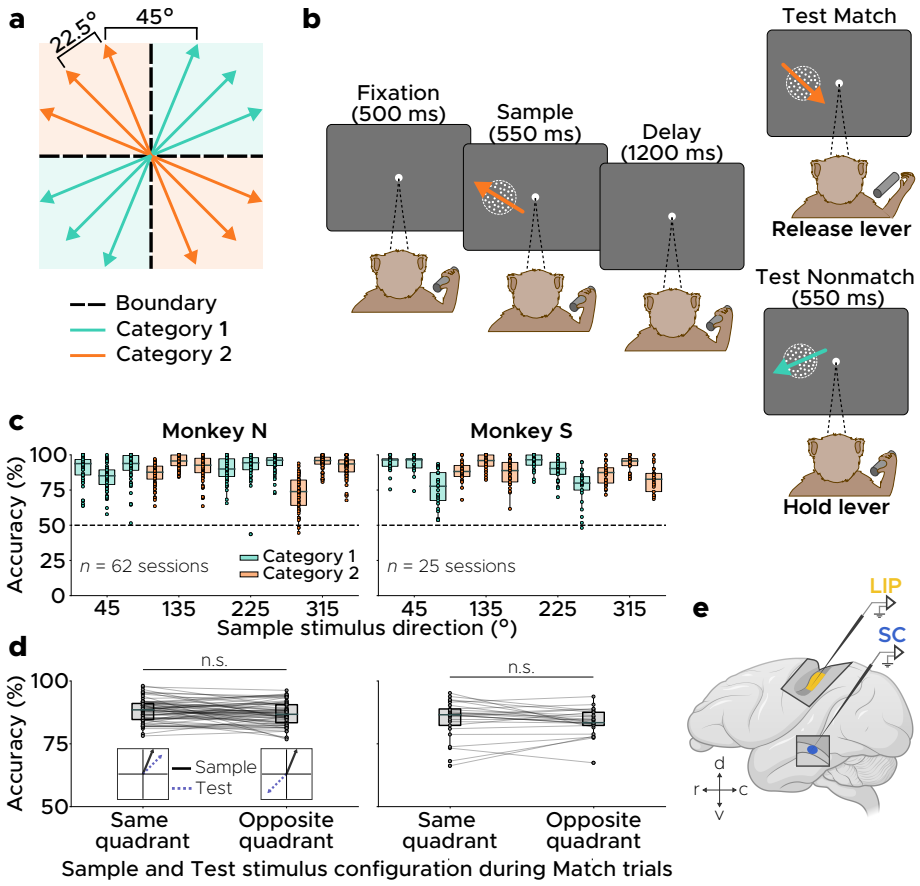

Figure 2.1: Monkeys learn to categorize motion stimuli based on an arbitrary category rule. a, Stimulus geometry of the two-boundary delayed match-to-category (DMC) task. 12 directions of motion are grouped into two categories based on two orthogonal category boundaries (dashed lines), such that directions that are 180° apart belong to the same category. Directions within the same quadrant are 22.5° apart, and near-boundary directions are 22.5° from the boundary. **b,** Trial structure of the DMC task. Monkeys were required to maintain gaze within a small window centered on a central fixation cue, and reported their decisions with a hand movement (holding or releasing a lever). **c,** Behavioral performance across recording sessions for each of the 12 sample stimulus directions for Monkey N (left) and Monkey S (right). Horizontal dashed line indicates chance performance. **d,** Behavioral performance across sessions on Match trials in which the sample and test stimuli were in the same or opposite quadrants. There was no significant difference in mean performance on trials in which sample and test stimuli were in the same or different quadrants (Monkey N: Same quad. = 88.2 ± 4.6%, Opp. quad. = 86.9 ± 4.70%, $P = .094$, Monkey S: Same quad. = 84.9 ± 7.4%, Opp. quad. = 83.8 ± 5.2%, $P = .580$, permutation test). **e**, Schematic of neural recording locations in the lateral intraparietal area (LIP) and superior colliculus (SC).

Following long-term training, both monkeys performed the DMC task with >85% mean accuracy during neural recording sessions (Monkey N: 89.40 ± 3.43%, Monkey S: 88.20 ± 3.34%; **Fig. 2.1c**), with no significant difference in mean performance between LIP and SC recording sessions (Monkey N: $P = .094$, Monkey S: $P = .229$, permutation test; **Extended Fig. 2.5**). Monkeys performed similarly on Match trials in which sample and test stimuli were in the same vs. opposite quadrants (Monkey N: $P = .094$, Monkey S: $P = .580$, permutation test; **Fig. 2.1d**), indicating that they learned to categorize stimuli across opposite quadrants.

2.2.2 Robust encoding of sample category in the SC

We analyzed spiking activity during the DMC task from 555 LIP neurons (Monkey N: 228, Monkey S: 327) and 604 SC neurons (Monkey N: 362, Monkey S: 242; **Fig. 2.1e**). Individual LIP neurons often showed binary-like category selectivity during the sample and delay task periods, with distinct activity for directions in different categories and similar activity for directions in the same category (**Fig. 2.2a**), consistent with previous studies that used a similar categorization task with a simpler linear category boundary (Freedman and Assad, 2006; Swaminathan and Freedman, 2012; Rishel et al., 2013; Sarma et al., 2016; Mohan et al., 2021). This category selectivity extended even to stimuli in opposite quadrants that belong to the same category. Remarkably, individual SC neurons also showed strong category selectivity during sample stimulus presentation and the subsequent delay and test periods (**Fig. 2.2b**).

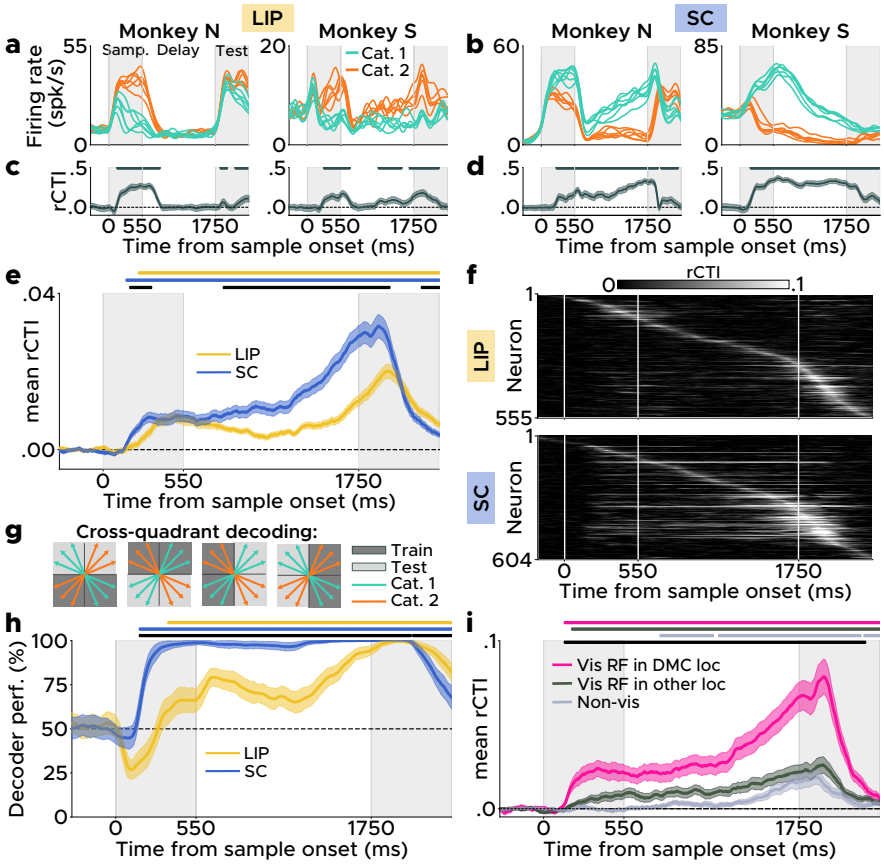

Figure 2.2: Neural activity in the SC contains reliable and short-latency sample stimulus category information. **a,** Peristimulus time histograms of example category-tuned LIP neurons. **b,** Same as *a*, but for SC. **c,** ROC-based category tuning index (rCTI) across time in trial for the two example LIP neurons in **a**. Shading indicates s.d., computed via resampling of trials. Grey symbols at the top of the plots indicate time points at which rCTI is significantly above chance. **d,** Same as *c*, but for SC. **e,** Time course of mean rCTI across LIP and SC neurons. Shading indicates s.e.m. **f,** Matrix of rCTI for all neurons in LIP and SC, where each row shows a single neuron's rCTI as a function of time in the trial. **g,** Schematic of cross-quadrant sample category classifiers, which are trained on trials from two quadrants (dark gray) and validated on trials from the remaining two quadrants (light gray). **h,** Time course of mean accuracy of cross-quadrant category classifiers for LIP and SC populations. Shading indicates s.d. **i,** Time course of mean rCTI across SC neurons that are visually responsive to stimuli at locations that overlap with the position of the DMC stimuli (pink), neurons that are visually unresponsive at DMC locations but visually responsive at other locations (dark grey), and neurons that are visually unresponsive (light grey). In **e**, **h** and **i**, colored symbols above plots indicate timepoints at which values significantly exceed chance. In **e** and **h**, black symbols above panel indicate time points at which there is a significant difference between brain areas, and in **i**, black symbols above panel indicate time points at which mean rCTI of Vis neurons is significantly higher than both Vis-other and Non-vis neurons ($P < .050$, two-tailed permutation tests).

We quantified the strength, trial-by-trial reliability, and time course of category tuning in individual neurons using an ROC-based category tuning index (rCTI), which compares neuronal discrimination between pairs of directions in the same vs. different categories (**Extended Fig. 2.6a**; see *Methods*). rCTI values can range from -0.5 to 0.5, with positive values indicating larger differences in firing rates between pairs of directions in different vs. same categories (and thus strong category tuning) and negative values indicating the opposite. **Fig. 2.2c** and **d** show the time course of rCTI for the single-neuron examples in **Fig. 2.2a** and **b**, and **Fig. 2.2f** shows a heatmap of rCTI values for all neurons in LIP and SC. We classified neurons as category-tuned if they showed significantly elevated rCTI values relative to a null rCTI distribution for at least 30 consecutive ms; **Extended Fig. 2.6**; see *Methods*).

In the LIP, 69.5% of neurons were category-tuned (Monkey N: 80.7%, Monkey S: 61.8%), and in the SC, 60.3% of neurons were category-tuned (Monkey N: 66.0%, Monkey S: 51.7%). In both LIP and SC, mean rCTI across neurons was shifted toward positive values at nearly every time point of the DMC task following sample onset (**Fig. 2.2e**). The increase in mean rCTI above baseline levels during the sample period of the task occurred significantly earlier in the SC than in the LIP (LIP: 245 ms, SC: 160 ms, $P = .008$, permutation test; **Fig. 2.2e**, yellow and blue symbols above panel). Moreover, rCTI values were significantly greater in the SC than the LIP throughout much of the trial (**Fig. 2.2e**, black symbols above panel), indicating stronger category encoding in the SC than LIP. During the sample epoch, a significantly greater percentage of LIP neurons were direction tuned compared to SC (LIP: 30.2%, SC: 20.4%; $X^2 = 7.65$, $P = .006$, X^2 test). To determine whether the earlier and stronger category tuning in the SC compared to LIP can be explained by the difference between brain areas in the proportion of direction-tuned neurons, we computed the mean rCTI in each area only on direction-untuned neurons (see *Methods*). This analysis again revealed shorter-latency category selectivity in SC compared to LIP (LIP: 195 ms, SC:

155 ms; $P = .042$, permutation test).

We next used linear support vector machine (SVM) classifiers to quantify the amount and timing of direction-independent category encoding in the LIP and SC neural populations. To evaluate the strength of category encoding in a direction-independent manner, the classifiers were trained on trials from two quadrants (one from each category) and validated on the remaining two quadrants (**Fig. 2.2g**). The logic behind this classifier is that, if the neural populations robustly encode category in a binary-like format, the classifier will be able to generalize between the two quadrants within the same category. In addition, this approach prevents direction tuning from contributing to category decoding performance by decorrelating direction and category between the sample and test sets; note that we find below-chance classifier performance when the population shows strong direction tuning. We also assessed the strength and time course of motion direction encoding in the neural populations from both brain areas using a category-independent direction classifier, as shown in **Extended Fig. 2.7**.

In the SC, category classifier accuracy rapidly increased above chance levels within approximately 170 ms of sample stimulus onset, and remained at almost 100% throughout the rest of the trial (**Fig. 2.2h**). In the LIP, category classifier accuracy was below chance shortly after sample stimulus onset, indicating that direction tuning was more dominant than category tuning during the early sample epoch. Sample category information could be decoded more reliably from SC than LIP activity throughout the sample, delay, and early test phases of the task (**Fig. 2.2h**, black symbols above panel), indicating stronger sample category encoding in the SC neural population than the LIP population.

Results regarding the strength and timing of category selectivity for the rCTI and decoding analyses were qualitatively similar in the two animals (**Extended Fig. 2.8**). The remarkably strong and short-latency encoding of abstract category in the SC is surprising, since primate SC is strongly associated with oculomotor control, orienting, and target selec-

tion, as opposed to higher-order non-spatial cognitive functions like categorization, and the animals here had to maintain central gaze fixation and did not report their decisions with saccadic eye movements.

2.2.3 Category encoding in the SC cannot be explained by eye movements

Given SC's well-established role in directing saccadic eye movements (Adamük, 1870; Apter, 1946; Wurtz and Goldberg, 1971; Schiller and Koerner, 1971; Schiller and Stryker, 1972; Robinson, 1972; Wurtz and Goldberg, 1972; Sparks, 1978; Hikosaka and Wurtz, 1985, 1986; Lee et al., 1988), one explanation for the unexpected presence of category selectivity in the SC is that the reported category signals could be a result of distinct patterns of microsaccades during different conditions of the DMC task. Interestingly, we observed that the monkeys produced idiosyncratic, category-specific eye movements (within the allowed fixation window) that were highly stereotyped across sessions (**Extended Fig. 2.9**). The category-specific microsaccades occurred only during the memory delay epoch in Monkey N and primarily during the memory delay epoch in Monkey S, indicating that the monkeys' eye movements reflect their working memory contents, consistent with results from past work in monkeys performing a delayed matching task (Dotson et al., 2018). To understand whether these eye movements can explain the observed category signals in the SC, we first compared the time course of neuronal category selectivity in the SC and the time course of category-specific eye position. This revealed that the two time courses were highly decoupled in time (**Extended Fig. 2.10a**), with category selectivity preceding category-specific eye position by several hundred milliseconds (Monkey N: neural decoder = 175 ms, eye decoder = 1075 ms; Monkey S: neural decoder = 170 ms, eye decoder = 865 ms). We next built linear encoding models (Musall et al., 2019) to determine whether firing rates of individual neurons (across trials and time within trial) are better predicted by stimulus category or eye movements (see *Methods*). In the majority of SC neurons, firing rates during the DMC

task were better predicted by the stimulus than by eye movements (**Extended Fig. 2.10f**). These results indicate that category selectivity in the SC cannot be accounted for by the category-specific eye movements during the DMC task, and raise the possibility that the eye movements may rather be a consequence of the presence of category selectivity in the SC.

2.2.4 Preferential encoding of category in visual SC neurons

The SC is a core stage of oculomotor processing and contains a diversity of neuronal response types based on firing rate modulation to visual, visuomotor, and motor aspects of visually- and memory-guided saccade (MGS) tasks. We were interested in understanding whether abstract category encoding in the SC was more prevalent among SC neurons with particular patterns of visual or motor selectivity. At the beginning of each DMC recording session, monkeys performed a block of the MGS task (**Extended Fig. 2.11**a,b), allowing a comparison of neuronal activity and selectivity from the same neurons during the two tasks. We analyzed activity from a subset of 423 SC neurons (Monkey N: 259, Money S: 164) from which we recorded during both the DMC and MGS tasks. Examples of single-neuron responses from SC during the MGS task are shown in **Extended Fig. 2.11c** and **d**).

The short-latency category encoding in the SC raises the possibility that the SC plays a direct role in the rapid bottom-up categorization of incoming visual stimuli (i.e., the transformation of direction information into category representations). One piece of evidence that would support such a role is if the category signal first emerges in visually-responsive neurons whose receptive fields match the position of the DMC stimuli. We therefore examined whether SC category selectivity varied between three groups of neurons: (1) those that are visually responsive during the MGS task to stimuli presented at locations that overlap with the stimulus position in the DMC task (Vis neurons), (2) neurons that are visually unresponsive at the DMC stimulus location but visually responsive at other locations (Vis-other), and (3) visually unresponsive neurons (Non-vis; see *Methods* for details). We compared mean

rCTI for 115 Vis neurons (Monkey N: 82, Monkey S: 33), 178 Vis-other neurons (Monkey N: 97, Monkey S: 81), and 130 Non-Vis SC neurons (Monkey N: 80, Monkey S: 50). Category selectivity emerges significantly earlier in Vis neurons than in Vis-other neurons (Vis: 145 ms, Vis-other: 195 ms, $P = .036$, permutation test) or Non-vis neurons (Non-vis: 800 ms, $P < .001$, permutation test; **Fig. 2.2i**), and mean rCTI is significantly higher in Vis neurons compared to the other two groups throughout much of the trial (**Fig. 2.2i**, black symbols above panel).

2.2.5 Orthogonal encoding of saccades and stimulus category in the SC

Next, we sought to understand the structure of population activity that allows motor and non-motor representations to coexist within a core oculomotor region like the SC. How is it that the SC can be strongly modulated by stimulus category (or visual information in general) without producing task-interfering saccades, given that injection of even a small amount of electrical current into intermediate and deep layers of the SC can reliably generate large-amplitude eye movements (Robinson, 1972; Schiller and Stryker, 1972)? One explanation is that independent populations of SC neurons might participate in saccade and category encoding. To investigate this possibility, we characterized the amount of overlap in neural populations that are category-selective during the DMC task and saccade direction-selective during the MGS task (see *Methods*). We observed substantial overlap in the population of SC neurons that are category-selective during the DMC task and saccade-direction-selective during the MGS task (**Fig. 2.3a**). In addition, we compared neuronal responses during the peri-saccade period of the MGS task and sample period of the DMC task, and calculated a task-preference index to quantify the ratio of how much each neuron is modulated by the two tasks (see *Methods*). If there are separate populations of neurons that encode saccade parameters and category, the distribution of task-preference indices across neurons would be bimodal. However, this was not observed in the data: the distributions peaked near

zero, indicating that a majority of neurons participated in both tasks (**Fig. 2.3a**), with no evidence for bimodality (Hartigan's dip test, Monkey N: $P = 0.984$, Monkey S: $P = 0.988$).

Another possibility is that the structure of population activity in the SC is organized to maintain approximately orthogonal neural representations during category processing and saccade planning, such that projection of category-related neural activity onto the saccade-encoding neural axis produces minimal eye movements (**Fig. 2.3c**). To investigate this idea, we quantified the degree of alignment between the neural activity space from the sample period of the DMC task (from +150 to +550 ms relative to sample onset) and the peri-saccade period of the MGS task (from -200 to 0 ms relative to saccade onset). This approach (Elsayed et al., 2016) compares the percentage of DMC data variance explained when the DMC data are projected onto DMC-defined vs. MGS-defined neural axes, and produces a single alignment index (AI) value that ranges from 0 (indicating perfect orthogonality between the two subspaces) and 1 (indicating perfect alignment). To determine whether the resulting indices are more or less aligned than random, we compared the AI computed from the data to a null distribution of AI values between subspaces drawn from a random space that shares a covariance structure with the real data (see *Methods*).

In both monkeys, alignment between the category and saccade subspaces was near-orthogonal; projection of the DMC data onto the MGS axes captured minimal DMC data variance (**Fig. 2.3d**), and the resulting alignment index was significantly closer to 0 than expected by chance (**Fig. 2.3e**; Monkey N: AI = .159, $P < .001$, Money S: AI = .222, $P < .001$). This result is unlikely due to neural fluctuations over time in the session; in both monkeys, DMC activity from the beginning of the sessions was closely aligned to DMC activity from end of the sessions (**Extended Fig. 2.12a**), and alignment to MGS activity was similarly low for DMC activity from the beginning vs. end of the sessions (**Extended Fig. 2.12b**).

To determine whether this misalignment between category and saccade subspaces may

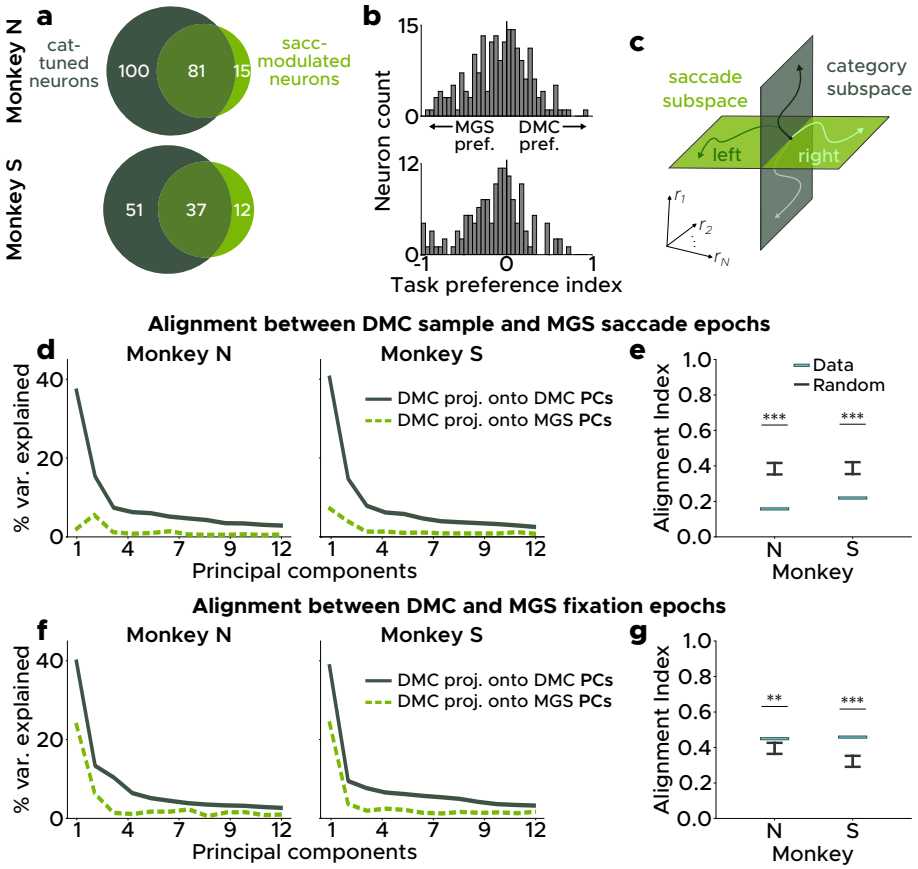

Figure 2.3: Orthogonal population-level encoding of saccade and category in the SC. a, Venn diagrams showing the overlap of neurons that are category-tuned during the DMC task (dark green) and neurons that are saccade-direction-modulated during the MGS task (light green; see *Methods*). **b,** Distributions of task preference indices across neurons for Monkey N (bottom) and Monkey S (top). Positive values indicate preference for the DMC task, and negative values indicate preference for the MGS task. **c,** Schematic of hypothesized orthogonal saccade and category activity subspaces. **d,** The percentage of variance of DMC sample period data explained when projected onto its own top 12 principal components (solid dark green line) or onto the top 12 principal components defined by activity during the peri-saccade period (-200 to 0 from saccade onset) of the MGS task (dotted light green line). **e**, Alignment index between the DMC sample epoch data and MGS peri-saccade data. The alignment index, which is the ratio of the two traces shown in **d**, equals 1 when two subspaces are perfectly aligned and equals 0 when two subspaces are perfectly orthogonal. Blue: alignment indices for the real data. Grey: 95% confidence intervals of alignment indices between pairs of random vector projections from data (see *Methods*). For both monkeys, the real data are significantly more orthogonal than expected by chance ($P < .001$). **f** and **g**, same as *d* and *e* but for the fixation epochs (-500 to 0 ms relative to stimulus onset) for the DMC and MGS tasks. The data are significantly more aligned than expected by chance (Monkey N: $P = .004$, Monkey S: $P < .001$).

be due to general differences in behavioral state in the two task contexts, we quantified the degree of alignment between the baseline neural activity during the fixation epochs (from -500 to 0 ms relative to stimulus onset) for the two tasks. During this baseline period, the task demands (i.e., maintaining fixation) are shared between the two tasks, but the overall behavioral context is different. In both monkeys, fixation epoch activity during the DMC and MGS tasks was significantly more aligned than expected by chance (**Fig. 2.3f-g**; Monkey N: AI = .451, $P = .007$, Money S: AI = .460, $P < .001$), indicating that the misalignment between category and saccade subspaces cannot be explained by differences in behavioral state between the two tasks, and suggesting that the SC may selectively use an orthogonal-coding strategy to minimize motor interference.

Together, these results suggest a mechanism by which neural populations in the SC can multiplex motor signals and the higher-order, non-motor cognitive signals that we report here. The SC may use a similar mechanism to encode visual information during contexts in which animals do not need to produce (or are required to withhold) eye movements. These results also provide a possible explanation for the stereotyped, category-specific microsaccades that emerge several hundred ms after neural category selectivity onset during the DMC task (**Extended Fig. 2.9**); these category-specific eye movements may reflect "leak" from the category subspace to the saccade subspace. During learning of the DMC task, the SC network may arrive at a solution (i.e., a particular geometry of population activity) that is sufficiently (though not perfectly) orthogonal to the saccade subspace, such that any resulting eye movements are within the behavioral constraints of the task (i.e., fall within the allowed fixation window).

2.2.6 Reversible pharmacological inactivation of the SC impairs performance on DMC task

Finally, we sought to determine whether category-correlated neural activity in the SC plays a causal role in the DMC task by infusing muscimol, a $GABA_A$ agonist, to reversibly inactivate the SC (**Fig. 2.4a**). We compared monkeys' behavior during SC inactivation with that during control sessions collected on the same day before injection.

We first verified the efficacy of SC inactivation by monitoring changes in saccade velocity for saccades towards the inactivated hemifield during a memory-guided (Monkey N) or visually-guided (Monkey S) saccade task. With the small dose of muscimol injected into SC (see **Extended Table 2.2**), both monkeys were able to successfully perform the memory/visually-guided saccade task following injection, even for targets in the inactivated hemifield. However, both monkeys exhibited a large reduction in accuracy and peak velocity of saccades to targets contralateral to the inactivated hemisphere (**Fig. 2.4b-d**), an effect that is characteristic of reversible SC inactivation (Hikosaka and Wurtz, 1985, 1986; Lee et al., 1988). This difference in mean peak velocity between the control and post-treatment blocks was significant for data combined across all muscimol injection sessions (Monkey N: control = 418 ± 90°/s, treatment = 249 ± 51°/s, $P < .001$; Monkey S: control = 420 ± 106°/s, treatment = 193 ± 92°/s, $P < .001$, permutation test; **Fig. 2.4d**), and was consistent (and significant) for all individual muscimol injection sessions (**Extended Table 2.3**). This effect was absent for data across sessions in which the SC was infused with saline (Monkey N: control = 432 ± 90°/s, treatment = 438 ± 95°/s, $P = .673$, Monkey S: control = 415 ± 84°/s, treatment μ= 413 ± 87°/s, $P = .848$, permutation test) and for all individual saline injection sessions (**Extended Table 2.3**).

Both monkeys showed a marked impairment in DMC task performance during muscimol injection sessions (**Fig. 2.4e-f**). We observed a significant reduction in accuracy on the DMC task in every experimental session in which SC was injected with muscimol, and no

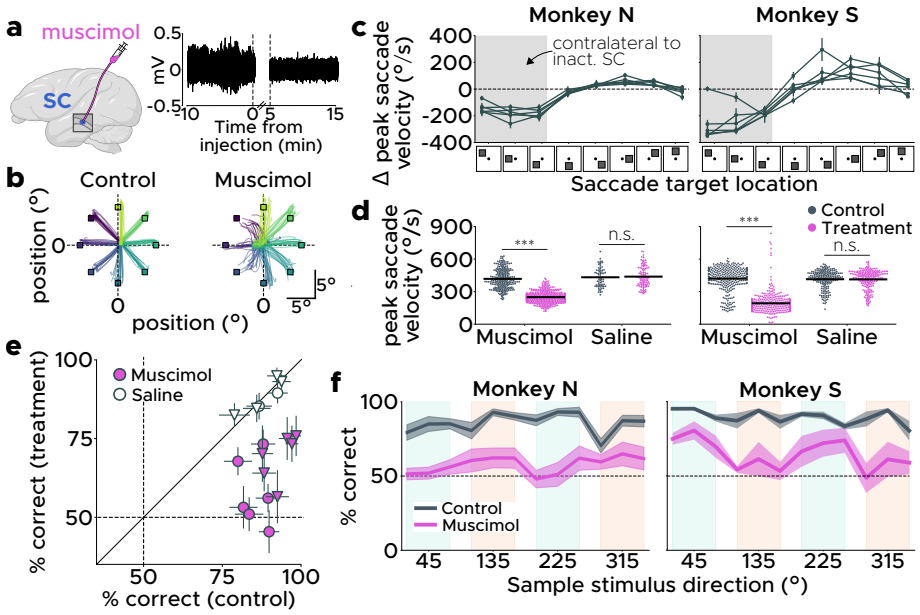

Figure 2.4: SC is causally involved in categorization task performance. a, We reversibly inactivated the SC using muscimol, a $GABA_A$ agonist (left). Shortly following injection of muscimol, spiking activity markedly decreased in the surrounding tissue (right). **b,** Trajectories of eye movements in Monkey S during the visually-guided saccade task for an example session before (left) and after (right) muscimol injection. Individual traces represent eye gaze trajectories for individual trials and are color-coded by condition. Small colored squares indicate the position of saccade targets for different conditions. **c,** Changes in peak saccade velocity during the memory/visually-guided saccade task before and after SC inactivation at each of the eight stimulus locations used in the task. Grey background shading indicates conditions in which the target location was in the inactivated hemifield. **d,** Peak saccade velocity for trials (across all sessions) in which the target was in the inactivated hemifield during control blocks (black) and treatment blocks (pink) for muscimol and saline injection sessions. Horizontal black lines indicate mean of distributions. **e,** Overall session accuracy for the DMC task before vs. during SC inactivation. Unfilled: saline injection, Filled: muscimol, Circles: Monkey N, Triangles: Monkey S. Error bars: 95% multinomial confidence intervals for each session. **f,** Mean performance on the DMC task on each of the 12 sample motion stimuli during control behavior (dark grey) and SC inactivation (pink). *** $P < .001$, permutation test.

change in accuracy for saline injection sessions (**Extended Table 2.4**). These results are consistent with SC being causally involved in performance of the DMC task.

The behavioral impairment on the DMC task during SC inactivation is unlikely to be entirely due to deficits in low-level visual processing of stimuli presented in the inactivated hemifield. The DMC deficit is not purely an attentional/hemispatial neglect-like impairment,

as the monkeys are still able to perceive, attend to, and respond to stimuli presented in the inactivated hemifield during the MGS task, and still attempt the DMC task and respond at appropriate times in the trial (i.e., release the lever only during the test epochs). The DMC deficit is also unlikely due to an impairment in the sensory processing of motion stimuli, as previous studies using similar experimental protocols have shown that SC inactivation produces only very slight impairments in direction discrimination of high-coherence motion stimuli like those used in our DMC task (Lovejoy and Krauzlis, 2010).

It is important to note that our experimental design cannot isolate the precise nature of the deficit caused by SC inactivation, as the DMC task requires several complex computations, including the transformation of sample direction into category, the maintenance of sample category information in working memory, the computation of the test stimulus category, and the comparison of sample and test stimulus categories. Future experiments with additional control tasks can more precisely characterize the nature of the deficit(s) caused by SC inactivation during this task. Despite this limitation, our study reveals that the functions of primate SC extend beyond those ascribed to it by current models.

2.3 Discussion

Our results demonstrate that the role of the primate SC extends beyond sensorimotor functions and spatial orienting to abstract, higher-order cognitive processing, even in a task that does not involve reporting decisions with saccades. The SC robustly encodes the learned categories of visual stimuli during all DMC task phases (including stimulus presentation, short-term memory, and comparison periods), and category signals in the SC arose with a shorter latency and were stronger than in LIP, an area previously shown to be causally involved in a similar categorization task. Moreover, reversible inactivation of the SC markedly impairs category task performance, indicating that activity in the SC is causally involved in abstract categorization.

Our findings also suggest that the SC neural population multiplexes category and saccade information by projecting those variables into distinct and near-orthogonal activity subspaces. This can explain how a motor structure like the SC can encode other task variables without producing task-interfering eye movements, closely related to the mechanism proposed to explain motor planning activity (without motor output) in the primary motor cortex (Elsayed et al., 2016). Even more broadly, this could be a general principle of neural coding through which a single neural population can efficiently and robustly encode multiple factors.

Previous studies that observed cognitive and/or abstract encoding in the SC used tasks in which animals either spatially orient to a particular target to indicate their choices (Horwitz and Newsome, 1999, 2001; Ratcliff et al., 2003; Horwitz et al., 2004; Ratcliff et al., 2007; Crapse et al., 2018; Jun et al., 2021; Stine et al., 2022; Duan et al., 2021, 2015; Felsen and Mainen, 2008), or tasks in which animals need to covertly orient to stimuli at distinct cued locations in different task conditions (Ignashchenkova et al., 2004; Cavanaugh and Wurtz, 2004; Müller et al., 2005; Cavanaugh et al., 2006; Lovejoy and Krauzlis, 2010; Zénon and Krauzlis, 2012; Krauzlis et al., 2013, 2014). By contrast, in the current study, monkeys did not report their decisions with a saccade or orient attention to different locations for different categories. Thus, it is difficult to account for our results based on differences in covert spatial attention. The DMC task also required monkeys to maintain gaze fixation throughout the trial, and category encoding in SC could not be explained by the animals' patterns of eye movements during the task.

We investigated SC during the DMC task because of evidence that cortical areas that are closely involved in oculomotor functions, such as LIP and the frontal eye fields (FEF), are also engaged in abstract categorization and flexible decision tasks (Ferrera et al., 2009). Our previous work shows that LIP plays a causal role in abstract categorization (Zhou and Freedman, 2019), and that it preferentially encodes motion categories compared both to

visual cortical areas MT and MST (Freedman and Assad, 2006; Zhou et al., 2022), and executive control areas such as the lateral prefrontal cortex (Swaminathan and Freedman, 2012). Anatomical connections between LIP and SC also motivated examining SC (Fries, 1984; Lynch et al., 1985; Asanuma et al., 1985; Andersen et al., 1990; Paré and Wurtz, 1997; Clower et al., 2001), as well as the similar patterns observed in SC and LIP during saccade-based tasks. We were also inspired by work showing that SC activity reflects higher-order functions such as attention and perceptual decisions (Jun et al., 2021; Duan et al., 2021; Crapse et al., 2018; Stine et al., 2022; Horwitz and Newsome, 1999, 2001; Ratcliff et al., 2003; Horwitz et al., 2004; Ratcliff et al., 2007; Ignashchenkova et al., 2004; Cavanaugh and Wurtz, 2004; Müller et al., 2005; Cavanaugh et al., 2006; Lovejoy and Krauzlis, 2010; Zénon and Krauzlis, 2012; Krauzlis et al., 2013, 2014). Our results highlight the need to directly and simultaneously compare encoding across the SC-FEF-LIP network to determine their contributions to computing abstract category information from sensory representations in upstream visual cortical regions (e.g. MT and MST) (Freedman and Assad, 2006; Zhou et al., 2022; Born and Bradley, 2005).

Monkeys were trained on the DMC task for hundreds of training sessions, spanning a period of many months, prior to neuronal recordings. It will be interesting to investigate whether abstract encoding in subcortical or motor structures such as the SC emerges only after prolonged training on a task. Task-related encoding might emerge at different rates or at different learning stages in the LIP and the SC; the SC may be recruited to participate in a cognitive task only once it is highly familiar. Alternatively, the SC might participate more broadly in learning and performing complex tasks and behaviors than previously appreciated, even during early stages of learning a task.

Our findings are interesting to consider from an evolutionary perspective, and highlight the importance of considering the functions of the SC between mammals and other vertebrates. In order to respond appropriately to stimuli, animals must rapidly combine sensory

encoding with more abstract, learned knowledge. While previous work in mammalian SC has emphasized its role in simple sensory-motor mapping, our work raises the prospect SC also mediates more flexible, and even cognitive, behaviors. This could be mediated by learning-dependent plasticity within SC and/or contextual/cognitive inputs from higher brain centers. Indeed, it could be advantageous for an area like the SC, which is close to both sensory input and motor output brain centers, to play such a role in order to facilitate rapid yet flexible behaviors. The idea that the SC is involved in mediating complex behaviors is especially plausible in non-mammalian vertebrate species that lack a neocortex, in which the tectum occupies a much larger fraction of brain volume and is known to play a major role in visual processing. Indeed, studies have found innate spatial encoding of stimulus size category in the optic tectum of untrained barn owls (Mysore and Knudsen, 2011; Mysore et al., 2011). In mammals and primates, this spatial orienting circuit may have evolved to rapidly compute more complex types of information (like the visual categories described here), while cortical pathways developed to allow for slower, but even more sophisticated processing.

2.4 Methods

2.4.1 Subjects

Two adult (13-15 years old) male rhesus macaques (Macaca mulatta) participated in the experiment (Monkey N: ∼12 kg, Monkey S: ∼13 kg). All procedures were in accordance with the University of Chicago Institutional Animal Care and Use Committee and the National Institutes of Health guidelines and policies.

2.4.2 Behavioral tasks

For the behavioral tasks described below, the monkeys were head restrained and seated in a primate chair inserted inside an isolation box (Crist Instrument), facing a 24-inch LCD

monitor on which stimuli were presented (1,920 x 1,080 resolution, refresh rate 60 Hz, 57 cm viewing distance). Reward delivery, stimulus presentation, behavioral signals, and task events were controlled by MonkeyLogic software (Asaad et al., 2013), running under MATLAB on a Windows-based PC. Gaze position was measured with an optical eye tracker (Eyelink 1000; SR Research, Ottowa, Canada) with a 1.0 kHz sample rate. For both tasks, monkeys initiated trials by holding a manual touch bar.

Delayed match-to-category task

We trained monkeys to perform a delayed match-to-category (DMC) task in which they grouped twelve dot-motion stimuli into two categories based on two orthogonal boundaries, such that motion directions that are 180° apart belong to the same category. Motion directions were separated into quadrants with three directions per quadrant, and stimuli within the same quadrant were 22.5° apart and near-boundary directions were 22.5° away from the boundary. The stimuli were 6°-diameter circular patches of white dots moving at a speed of 10°/s with 100% coherence, presented at 6.5-7.5° eccentricity in the contralateral visual field. Animals were required to fixate within a 2.5-3.5° window.

Memory-guided saccade task

We used a memory-guided saccade (MGS) task to identify visual and motor receptive fields of LIP and SC neurons (**Extended Fig. 2.11**). At the start of a trial, monkeys had to maintain fixation on the central cue for 500 ms, after which a white square target briefly appeared for 300 ms at one of eight peripheral locations (equally spaced and concentric at 6.5° eccentricity; see **Extended Fig. 2.11b**). The target presentation was followed by a 1000-ms delay period, after which the fixation cue disappeared and monkeys had to saccade to the remembered location of the visual target presented earlier in the trial.

2.4.3 Surgical procedures and electrophysiological recordings

Monkeys were implanted with a titanium headpost and a single recording chamber positioned over LIP and SC. Stereotaxic coordinates for chamber placement were determined from magnetic resonance imaging (MRI) scans obtained before implantation of recording chambers. LIP and SC recordings were conducted in separate sessions, typically using 16- and 24-channel linear Plexon V-probes (in which channels span 1.5-2.0 mm of tissue), a dura-piercing guide tube, and a NAN microdrive system (NAN Instruments). A small subset of recording sessions from one monkey were conducted using single epoxy-insulated tungsten electrodes (FHC, Inc). We used anatomical landmarks and responses during the MGS task to guide recordings. For SC recordings, we primarily targeted neurons in superficial and intermediate layers, although we also recorded neurons in deep layers as well due to the ~2mm span of recording channels on our probes. Neurophysiological signals were amplified, digitized, and stored for offline spike sorting (Plexon) to verify the quality and stability of neuronal isolation.

2.4.4 SC inactivation

We infused muscimol, a $GABA_A$ agonist, to unilaterally inactivate the SC. We built a microfluidic injectrode system to deliver small amounts of the drug or saline (0.25-0.5 µL; see **Extended Table 2.2**) using the protocol developed by (Vanegas et al., 2019). To ensure that we precisely injected the drug into superficial and intermediate layers of the SC, we used a custom 16-channel Plexon S-probe with a fluid delivery channel that allowed us to monitor neural activity during probe lowering and before injection. Before drug injection on each session, monkeys first completed a control behavioral session in which they performed at least 200 correct trials of the DMC task and at least 100 correct trials of the MGS task (Monkey N) or visually-guided saccade task (Monkey S). After monkeys completed the control behavioral session, we infused the drug and waited 15-25 minutes to begin the post-

treatment behavioral session. To verify success of SC inactivation, we compared saccade metrics (peak saccade velocity) during the MGS/VGS tasks during the control and post-treatment trials. We analyzed data from 12 muscimol injections sessions (Monkey N: 6 sessions, Monkey S: 6 sessions) and six control saline injection sessions (Monkey N: 2 sessions, Monkey S: 4 sessions). **Extended Table 2.2** provides information for each injection session, including muscimol and saline concentration, injection volume, and number of completed DMC trials.

2.4.5 Behavioral inclusion criteria

For electrophysiological recordings and inactivation experiments, we included sessions in which behavioral performance on each category for the DMC task was at least 75% (criterion applied only to control blocks for the inactivation experiments). We excluded six LIP recording sessions (four in Monkey N and two in Monkey S) from analyses due to poor behavioral performance. For the inactivation experiments, we excluded two sessions in Monkey S (one saline injection session with 66% accuracy for category 2 during the control block, and one muscimol injection session with 53% accuracy for category 2 during the control block).

For DMC analyses, we included well-isolated neurons for which we had data recorded during at least five correct trials for each sample direction. We analyzed spiking data during the DMC task from 555 LIP neurons recorded over 49 recording sessions (Monkey N: N neurons = 228, N sessions = 36; Monkey S: N neurons = 327, N sessions = 13) and 604 SC neurons recorded over 38 recording sessions (Monkey N: N neurons = 362, N sessions = 26; Monkey S: N neurons = 242, N sessions = 12). We collected and analyzed spiking activity during the MGS task in a subset of 423 SC neurons (Monkey N: $N = 259$, Monkey S: $N = 164$) for which we recorded data from least two correct trials for each MGS condition.

2.4.6 Data analysis

All analyses were performed in Python v3.7.3. Behavioral analyses for the DMC task (including those for inactivation) were performed on all completed trials (i.e., correct trials, misses on match trials, and false alarms on non-match trials). Unless otherwise specified, all neural analyses for the DMC task were performed only on correct trials. Behavioral and neural analyses for the MGS/VGS tasks were performed only on correct (completed) trials. All *P*-values are two-tailed unless otherwise specified. For neural analyses, spike trains for each neuron were smoothed using Gaussian kernel (σ = 30 ms). Eye tracker gaze position data were low-pass filtered to reduce noise using a 2nd-order Butterworth filter with a 70-Hz cutoff.

Behavioral performance

To compare differences in mean behavioral accuracy (across all sample directions) between LIP and SC recording sessions in each monkey (**Extended Fig. 2.5a**), we used a permutation test in which we randomly permuted mean accuracy values between the two brain areas (while preserving the number of sessions per area). We repeated this procedure for 5000 unique iterations to generate a null distribution of accuracy differences. To compare differences in behavioral performance on match trials in which the sample and test stimuli were in the same vs. opposite quadrants, we computed the difference in mean accuracy for same vs. opposite quadrant match trials for each session and used a permutation test (with 5000 iterations) to randomly permute the per-session accuracy values between the two conditions.

Quantifying single-neuron category tuning

We quantified the strength, reliability, and time course of single-neuron category tuning using a receiver operating characteristic (ROC)-based category tuning index (rCTI)(Rishel

et al., 2013). For each neuron, we applied ROC analysis to distributions of trial-by-trial firing rates and compared area under the ROC curve (AUC) values for eight pairs of sample motion directions that are in the same category (Within-Category; WC) and eight pairs of directions that are in different categories (Between-Category; BC). To ensure equalized angle differences between WC and BC pairs (and thus prevent direction tuning from contaminating rCTI), the WC and BC groups each included four direction pairs spaced 45° apart and four direction pairs spaced 135° apart (**Extended Fig. 2.6a**). We quantified rCTI at each timepoint as the mean rectified WC AUC subtracted from the mean rectified BC AUC:

$$\begin{aligned} rCTI = \tfrac{1}{8}\sum_{p=1}^{8} 0.5 + \Big|0.5 - AUC(BC_{p1}, BC_{p2})\Big| \\ - \tfrac{1}{8}\sum_{p=1}^{8} 0.5 + \Big|0.5 - AUC(WC_{p1}, WC_{p2})\Big| \end{aligned} \tag{2.1}$$

where BC_{p1} and BC_{p2} are the two directions in the p^{th} BC pair, and WC_{p1} and WC_{p2} are the two directions in the p^{th} WC pair.

We applied the rCTI analysis to smoothed spike trains (see above) across 5-ms time steps in the trial. To generate the error shading shown in **Fig. 2.2c** and **d**, we calculated rCTI for each neuron over 500 bootstraps using 15 trials per sample motion direction (with replacement). We generated null distributions of rCTI values for each neuron using a bootstrap analysis (repeated 5000 times) in which we randomly assigned (with replacement) eight direction pairs (four 45°-spaced and four 135°-spaced pairs) to each of the shuffled BC and WC groups (**Extended Fig. 2.6**). We defined category-tuned "runs" as time bins at which rCTI values significantly exceed the null distribution for a minimum of six consecutive analysis time bins (30 ms). We considered neurons to be category-tuned if they had at least one significant run, and defined latency of category selectivity for each category-tuned neuron as the first time bin of the earliest significant run.

To test for significant above-chance mean rCTI in each brain area (as shown in **Fig. 2.3f**), we used a permutation procedure in which we computed a null mean rCTI across neurons for each WC/BC-label-shuffling iteration. To test for a significant difference between brain areas in the onset time of category selectivity for mean rCTI, we compared the observed between-area latency difference to a null distribution of latency differences. For each of 5000 iterations, we randomly permuted neurons between the two brain areas (while preserving the number of neurons in each area) and we computed the difference in latency of category selectivity onset in the two shuffled groups. To test for differences in onset time of category selectivity between Vis and Vis-other SC neurons, we used a similar procedure in which we randomly permuted neurons between the two groups instead of between brain areas.

Support vector machine (SVM) analyses

We used SVM classifiers (with a linear kernel) to quantify the strength and timing of sample stimulus category encoding in populations of LIP and SC neurons. To quantify category encoding in a direction-independent manner, we constructed cross-quadrant classifiers for which training sets consisted of trials in which the sample motion directions were from two of the four quadrants (one from each category), and testing sets consisted of sample motion direction trials from the other two quadrants (**Fig. 2.2g**). The training and testing quadrants were randomly chosen on each iteration. The analysis was applied in 5-ms steps across time in the trial and repeated for 200 iterations. For each neuron, we included 15 trials from each of six sample motion directions for training (as described above) and 15 trials from each of the remaining six sample motion directions for testing. To reduce the biases in classifier performance across brain areas due to an unequal number of neurons, for each iteration of the analysis, we randomly selected N neurons for inclusion, where N is the number of neurons in the brain area with the lower number of neurons. We generated null distributions of decoder performance values at each time using a permutation procedure (repeated 1000

times) in which we shuffled the sample direction label assigned to each trial.

We also used linear SVM classifiers to decode sample direction from LIP and SC population activity. To quantify the amount of direction encoding in a category-independent manner, the training/validation sets for each iteration of the classifier only included data from one of the two categories. The classifiers were trained on 48 trials (8 trials from each of the six directions from one of the two categories, randomly chosen) and validated on 12 held-out trials (2 trials from each of the six motion directions). This analysis was applied in 5-ms steps across the trial and repeated for 200 iterations.

Identifying task-responsive neurons

To identify neurons that are task-responsive during the DMC task, we used a bin- and parameter-free statistical test to detect any consistent time-locked modulations in firing rate for each neuron (Montijn et al., 2021). In brief, this analysis consists of the following steps (applied separately for each sample direction): (1) aligning the spike trains for all correct trials to the onset of the sample stimulus, (2) stacking these spike train to create a single vector of spikes relative to sample onset, (3) calculating the cumulative distribution of spikes over trial time using this spike vector, and (4) comparing this cumulative distribution to a linear baseline (which represents an unvarying firing rate over time), producing a deviation value for each timepoint. To generate a null distribution of 5000 deviation-from-baseline values, we shuffled the spike trains in each trial to destroy any time-locked activity patterns across trials while preserving the total number of spikes, and computed the maximum deviation (across time) for these shuffled data. For this analysis, we included data for each trial from 500 ms before sample stimulus onset (i.e., the start of the pre-sample fixation period) until the end of the first test stimulus epoch. We also computed the peak mean firing across time (in the period from the beginning of the sample epoch until the end of the first test epoch) for each sample direction. We classified each neuron as task-responsive if it satisfied

the following two criteria: (1) if it showed a significant modulation in firing rate (i.e., had significantly elevated deviation-from-linear-baseline values) at any timepoint from the start of the sample epoch until the end of the test epoch, and (2) if its maximum peak mean firing rate (across sample directions) was at least 3 sp/s. In LIP, 506/555 (91.2%) of neurons were task-responsive (Monkey N: 210/228, 92.1%; Monkey S: 296/327, 90.5%), and in SC, 504/604 (83.4%) of neurons were task-responsive (Monkey N: 300/362, 82.9%; Monkey S: 204/242, 84.3%).

Identifying direction-tuned neurons during the DMC task

To identify neurons that are significantly direction tuned during the sample epoch of the DMC task, we computed a direction tuning index (DTI) for each neuron using the circular variance method introduced in (Mazurek et al., 2014). We calculated neurons' mean firing rate for each sample stimulus direction in a direction vector space, and quantified DTI as the normalized length of the sum of these vectors:

$$DTI = \left| \frac{\sum_{k=1}^{12} f(\theta_k) e^{i\theta_k}}{\sum_{k=1}^{12} f(\theta_k)} \right| \tag{2.2}$$

where f(θ_{k}) is a neuron's mean firing rate for direction θ_{k}.

To test for significant direction tuning , we compared the true DTI to a distribution of 5000 null DTIs generated by randomly shuffling the direction labels assigned to each mean firing rate. We applied this analysis to firing rates in three non-overlapping 200-ms windows from 0 to +600 ms relative to sample stimulus onset. We classified neurons as direction tuned if they showed significant direction tuning during at least one of the three time windows and if they were identified as responsive during the DMC task (see above). In LIP, 153/506 (30.2%) of neurons were significantly direction-tuned (Monkey N: 79/210, 39.5%; Monkey

S: 74/296, 25.0%), and in SC, 103/504 (20.4%) of neurons were significantly direction-tuned (Monkey N: 53/300, 17.7%; Monkey S: 50/204, 24.5%).

Identifying visually-responsive neurons during the MGS task

We analyzed neuronal activity during the MGS task in order to characterize the visual and motor response fields of SC neurons. For each neuron, we determined whether the DMC stimulus was presented in its visual receptive field by identifying the MGS condition ($\text{MGS}_{\text{DMCloc}}$) whose location overlapped with the DMC stimulus location on that session. We then determined whether the neuron was significantly modulated during the MGS visual epoch for that condition. For each of the eight MGS conditions, we computed the mean firing rate (per trial) for each non-overlapping 25-ms bin from 0 ms to +400 ms relative to stimulus onset. We used the Kruskal–Wallis H-test to compare the firing rate distributions across these windows, and compared the resulting H-statistic to a distribution of 5000 null H-statistics. To generate the null H distribution, we shuffled the neuron's time-varying firing rates (from 0 to +400 ms relative to stimulus onset) for each trial, calculated the mean firing in non-overlapping 25-ms windows for these permuted trials, and computed a shuffled H-statistic. Neurons were classified as "Vis" neurons if they was significantly modulated across the visual stimulus period for the $\text{MGS}_{\text{DMCloc}}$ condition and if their maximum firing rate across analysis windows and conditions was above 3 sp/s. Neurons were classified as "Vis-other" if they were significantly modulated across the visual stimulus period for another MGS condition (and if their maximum firing rate was above 3 sp/s). Neurons were classified as "Non-vis" if they were not significantly modulated across the visual period for any of the MGS conditions, or if their maximum firing rate across analysis windows and conditions was below 3 sp/s. 115 (27.2%) neurons were classified as Vis neurons (Monkey N: 82 [31.7%], Monkey S: 33 [20.1%]), 178 (42.1%) as Vis-other neurons (Monkey N: 97 [37.5%], Monkey S: 81 [49.4%]), and 130 (30.7%) as non-Vis neurons (Monkey N: 80 [30.9%], Monkey S: 50

[30.5%]).

Identifying saccade-modulated neurons during the MGS task

We analyzed activity of SC neurons during the saccade period of the MGS task (-200 ms to +50 ms relative to saccade onset) to identify neurons that are significantly modulated by saccade direction. For each neuron, we computed its mean firing rate across the saccade period window for each trial. We used the Kruskal–Wallis H-test to compare the neuron's the firing rate distributions across the eight MGS conditions, and compared the resulting H-statistic to a distribution of 5000 null H-statistics generated by shuffling condition labels among trials. 145/423 (34.3%) neurons were significantly modulated by saccade direction (Monkey N: 96/259 [37.1%], Monkey S: 49/164 [29.9%]).

Task modulation index

To quantify differences in modulation during the MGS and DMC tasks for each SC neuron, we computed a task-preference index (**Fig. 2.3b**). We defined amount of MGS modulation as the range of mean firing rates across conditions during the peri-saccade period of the MGS task (-200 to +50 relative to saccade onset), and the amount of DMC modulation as the range of mean firing rates across sample directions during the sample epoch of the DMC task (+150 to + 550ms relative to stimulus onset). We then normalized these ranges for each neuron by the neuron's overall firing rates across all times/conditions/tasks. The task-preference index was defined as the ratio of DMC modulation to MGS modulation. We used a Hartigan's dip test to test for bimodality in the distribution of task-preference indices across all SC neurons.

Linear encoding models

We constructed linear encoding models (Musall et al., 2019) to quantify how much firing rates of individual neurons (across trials and time within trials) are modulated by stimulus category vs. microsaccades. The linear models contained regressors related to stimulus category and microsaccade parameters. For the category regressors, we constructed a binary vector containing a pulse at the time of the sample stimulus onset, and created copies of this vector shifted in time by 1ms for every point until the end of the trial. The microsaccade regressors included two analog regressors: horizontal and vertical eye velocity at each timepoint throughout the trial, shifted in time by -50 relative to neural activity to account for lag between neural activity and saccades. We also included two types saccade event kernel regressors: (1) a binary vector containing a pulse at every timepoint at which a microsaccade occurred, and (2) a vector containing microsaccade direction at every timepoint at which a microsaccade occurred and zeros at every other timepoint. We created time-shifted copies of the binary saccade vector and saccade direction vector, spanning from -500 until +100 ms (relative to saccade onset) in 10-ms steps. The design matrix of the full model included all of the category and saccade regressors. We also built reduced models that contained shuffled saccade regressors and unshuffled category regressors, or shuffled category regressors and unshuffled saccade regressors. For each neuron, we fit the models using ridge regression (with L2 regularization and 10-fold cross validation) and computed an R^2 for the full model and each of the reduced models. To quantify how well category or saccade regressors predict neural activity in each neuron, we computed the change in cross-validated R^2 from the full model to each reduced model. A large (negative) change in R^2 indicates a strong contribution of the excluded variables.

Subspace alignment analysis

We used a subspace alignment analysis introduced in (Elsayed et al., 2016) to quantify the degree of alignment between neural activity in the SC during the MGS and DMC tasks. For this analysis, we constructed matrices D and M of neural activity during the DMC and MGS tasks, respectively. D and M were size N by cxt, where N is the number of neurons, c the number of conditions (12 for DMC and 8 for MGS), and t is the number of time points per condition. Each row of D and M contains the concatenated mean firing rates (per condition and across time points) of one neuron. We normalized the firing rates of each neuron by its range (across all included DMC and MGS conditions and time points) plus a constant, chosen as 10 sp/s. We then performed principal components analysis (PCA) on the matrix D to obtain the top 12 DMC PCs, and on matrix M to obtain the top 12 MGS PCs. We then projected the DMC activity D onto both the DMC and MGS PCs and calculated sum of the percent of variance explained (relative to total variance of D) for each of the projections. We quantified the alignment index (AI) between the two subspaces as the ratio of these two sums. The logic behind this analysis is that if the DMC and MGS subspaces are approximately orthogonal, the projection of D onto the MGS PCs will capture minimal D variance. AI ranges from 0 (indicating perfect orthogonality between two subspaces) and 1 (indicating perfect alignment).

To determine whether measured AI values are more (or less) misaligned than expected by chance, we calculated the alignment between pairs of random subspaces sampled from the full covariance structure of the data to generate a null distribution of alignment values (Elsayed et al., 2016). To create the random subspaces, we first computed the covariance matrix C from the concatenated D and M matrices, and obtained the left singular vectors (U) and singular values (s) of C using singular value decomposition. For each of 5000 iterations per comparison, we computed the AI between two random subspaces (v_{rand}). We sampled each

random subspace v_{rand} as follows:

$$v_{rand} = orth\left(\frac{U\sqrt{S}\mathbf{v}}{||U\sqrt{S}\mathbf{v}||_2}\right) \tag{2.3}$$

where $\mathbf{v}$ is an N x 12 matrix in which each element is drawn from a normal distribution with mean 0 and variance 1, and $orth(X)$ returns the orthonormal basis of X defined by its left singular values.

For the main alignment analysis (shown in **Fig. 2.3c-d**), we included DMC task data from 150-550 ms after sample stimulus onset (during stimulus presentation) and MGS task data from 200 before saccade onset until saccade onset. Data for both tasks were sampled in 10-ms steps. Note that results were equivalent when we used time windows of equal length for the two tasks (MGS: -200 ms to +0 ms relative to saccade onset, DMC: +350 to +550 ms relative to sample onset; Monkey N: AI = .134, $P < .001$, Money S: AI = .233, $P < .001$).

SC inactivation analyses

To verify efficacy of SC inactivation, we quantified the difference in peak saccade velocity for saccades towards the inactivated hemifield during the MGS/VGS task between the control and post-treatment blocks. For each trial, we computed the maximum eye gaze velocity from 200 ms before go cue onset until successful target fixation initiation. We excluded one trial for Monkey N (session 2, muscimol treatment, upper-center condition) in which we could not accurately quantify peak saccade velocity because the monkey blinked during the response period. For each session, we combined trials from the three conditions in which the target was in the inactivated hemifield ("Contralateral"), and the three conditions in which the target was out of the inactivated hemifield ("Ipsilateral"). We tested for significant differences in mean peak saccade velocity between the control and treatment blocks for Contralateral and

Ipsilateral trials on each session using a bootstrap test with 5000 iterations (**Extended Table 2.3**). We also tested for significant differences in mean peak saccade velocity between control and treatment blocks for Contralateral trials pooled across all muscimol sessions and pooled across all saline sessions, as shown in **Fig. 2.4d**. For the muscimol sessions, this analysis included 270 (470) control (treatment) trials in Monkey N and 320 (376) control (treatment) trials in Monkey S, and for saline sessions included 72 (103) control (treatment) trials in Monkey N and 207 (215) control (treatment) trials in Monkey S.

We used a two-sided Fisher exact test to quantify differences in behavioral performance on the DMC task between control and post-treatment blocks for each session. The statistics for the test are shown in **Extended Table 2.4**.

2.5 Supplemental Figures and Tables

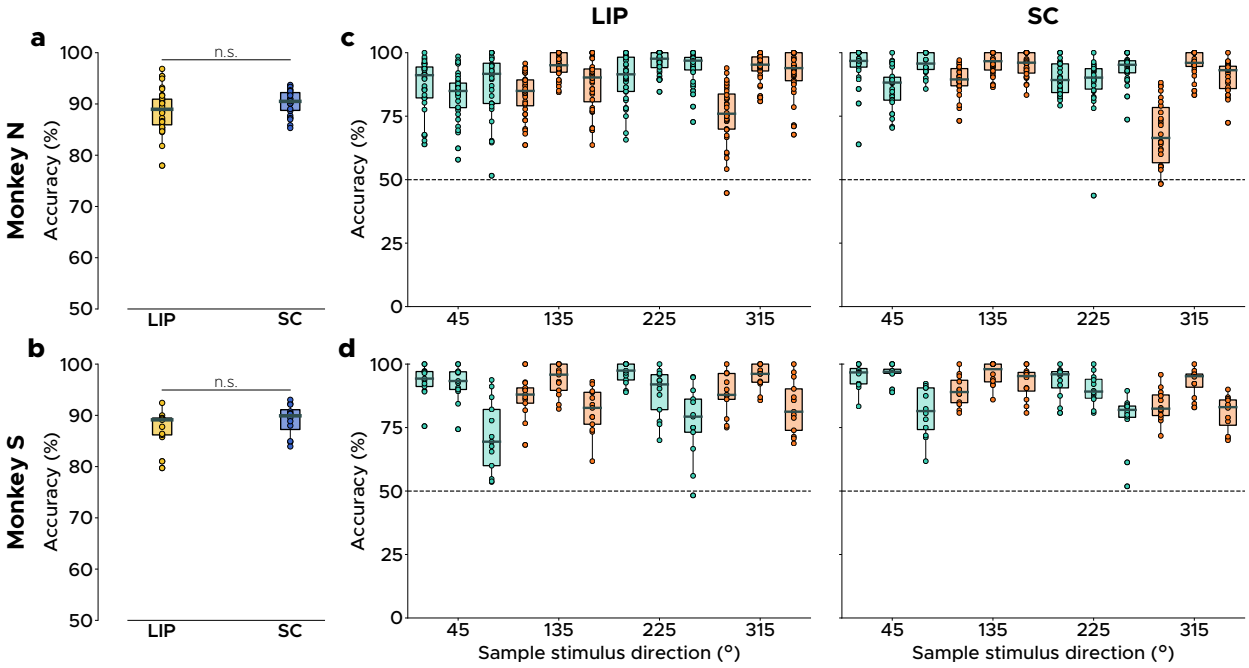

Figure 2.5: No difference in behavioral performance between LIP and SC recording sessions a, Distributions of overall session accuracy for LIP and SC recording sessions in Monkey N. There was no difference in mean accuracy between brain areas (LIP: 88.8 ± 3.9%, SC: 90.3 ± 2.3%, $P = .094$, permutation test). **b,** same as *a* but for Monkey S (LIP: 87.4 ± 3.5%, SC: 89.1 ± 2.9%, $P = .229$, permutation test). **c,** Mean accuracy by sample direction for LIP (left) and SC (right) recording sessions. **d,** same as *c* but for Monkey S.

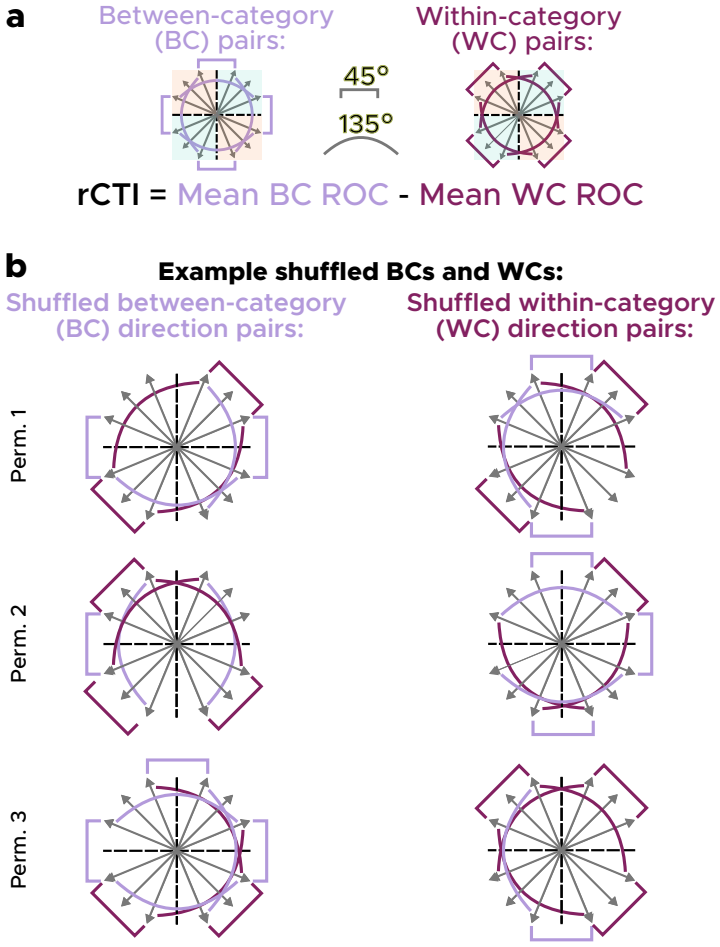

Figure 2.6: rCTI method and shuffling procedure a, Schematic of the rCTI analysis used to quantify strength of category tuning in each neuron. rCTI compares ROC values between pairs of sample directions that are in the same category (Within-Category; WC) vs. different categories (Between-Category; BC). WC and BC groups each consisted of eight sample directions (four pairs spaced 45° apart and four pairs spaced 135° apart. **b,** We generated null distributions of rCTI values across timepoints for each neuron using a shuffling procedure in which we reshuffled the labels (between- vs. within-category) assigned to each pair of directions, such that each shuffled group contained four 45°-apart direction pairs and four 135°-apart direction pairs. The permutation procedure was repeated 4900 times, once for every combination of shuffled direction pairs that conformed to the criterion above. Bracket colors denote true group assignment of a direction pair (light purple = between-category, dark purple = within-category).

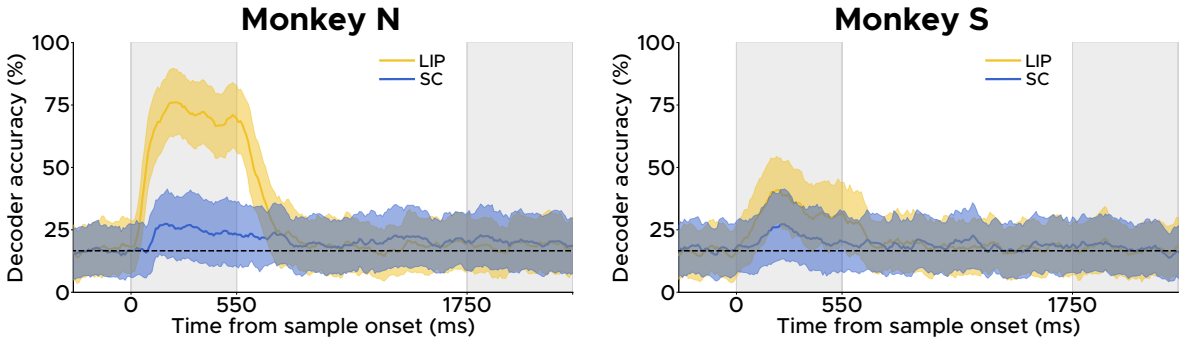

Figure 2.7: Direction classifier accuracy in LIP and SC Time course of category-independent direction classifier accuracy across LIP and SC neurons. Lines and shading indicate mean ± s.d.

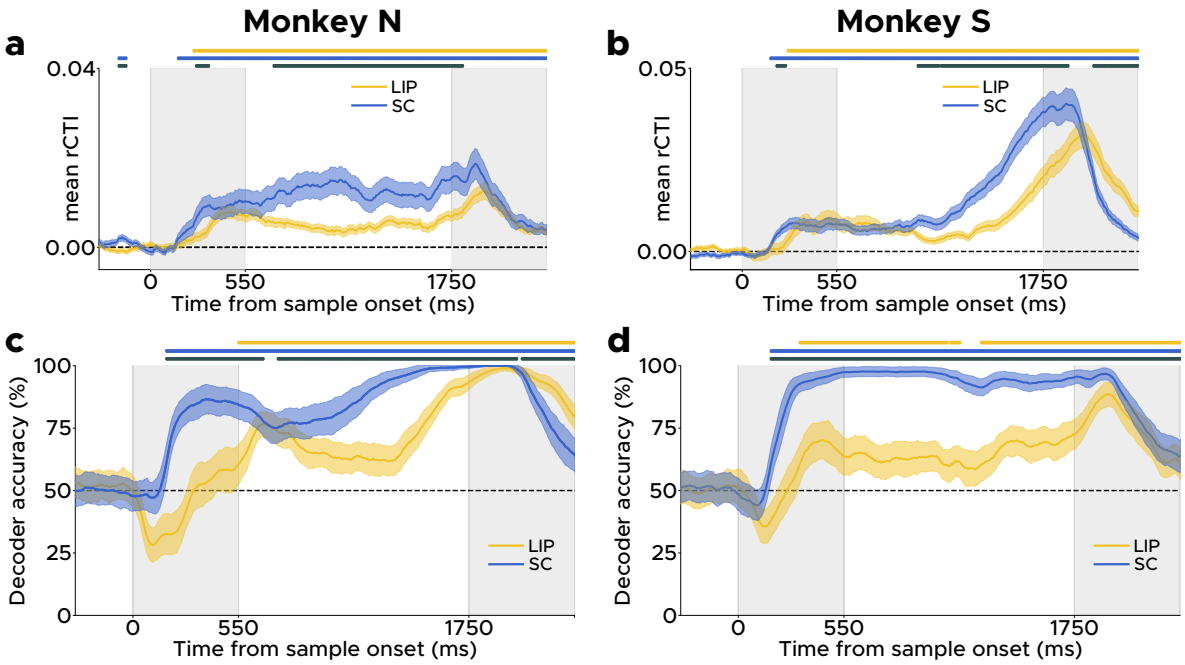

Figure 2.8: rCTI and category classifier accuracy in LIP and SC by monkey **a,** Mean Time course of rCTI across LIP and SC neurons in Monkey N. Lines and shading indicate mean ± s.d. rCTI. **b,** same as *a* but for Monkey S. **c,** Time course of cross-quadrant category decoding accuracy in LIP and SC for Monkey N. Lines and shading indicate mean ± s.d. decoding accuracy. **d,** same as *c* but for Monkey S. In all panels, yellow (blue) symbols above plot indicate timepoints at which LIP (SC) values are significantly above chance, and black symbols indicate time points at which there is a significant difference between LIP and SC values when either or both area are significantly above chance (permutation test, all $P <$ .050).

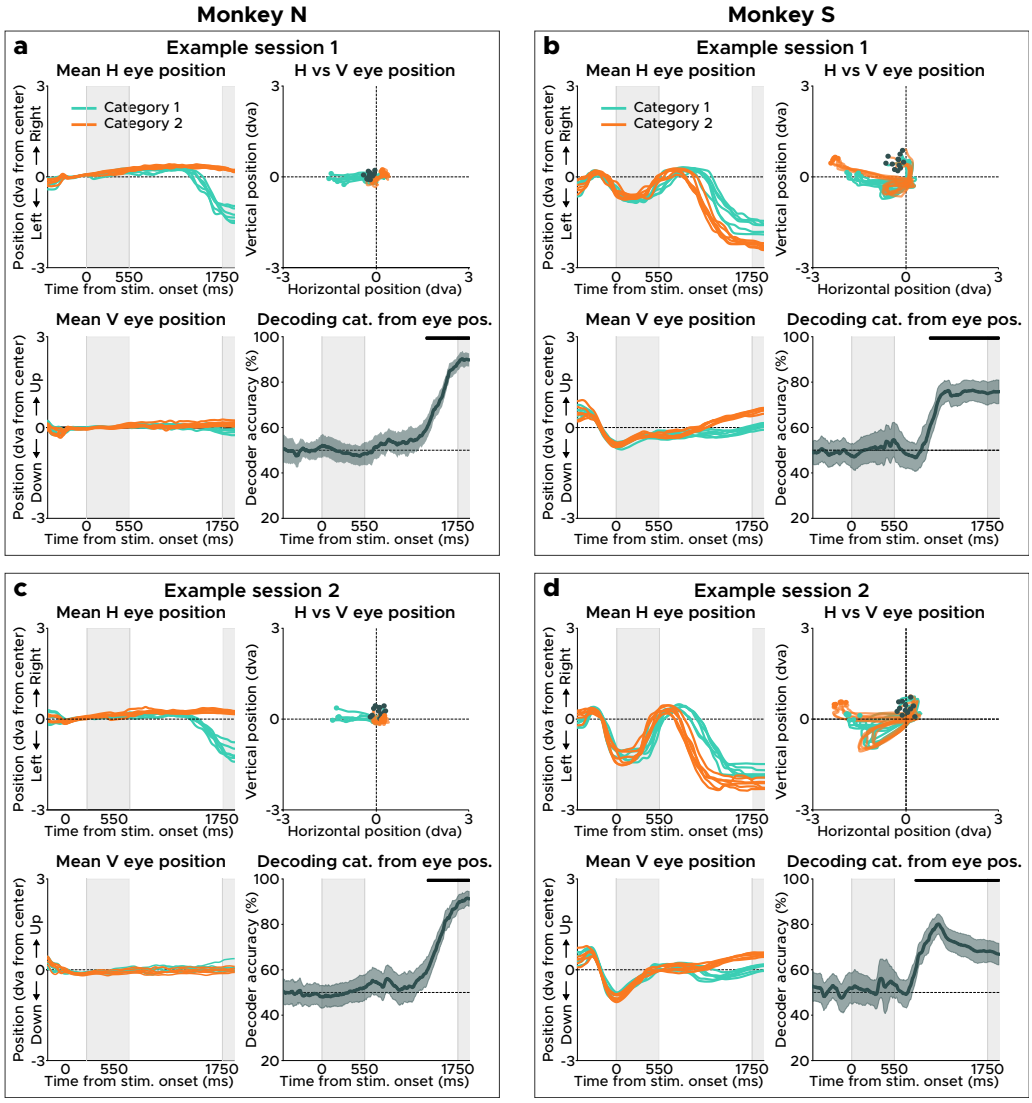

Figure 2.9: Monkeys' eye movements reflect working memory contents during delay period of DMC task a, Top left: Mean horizontal eye position across trial time for each sample motion direction during an example session. Bottom left: Mean vertical eye position across trial time. Top right: Mean horizontal vs. vertical eye position. Black circles indicate mean eye position at the begining of the trial, and colored circles indicate the mean eye position at the beginning of the delay period. Bottom right: Accuracy of cross-quadrant category classifier trained on eye position. Black symbols above panel indicate time points at which the category classifier performs significantly above chance (P <0.05, permutation test). Shading indicate s.d. **b,** Same as a but for Monkey S. **c, d** Same as a and b but for another example session.

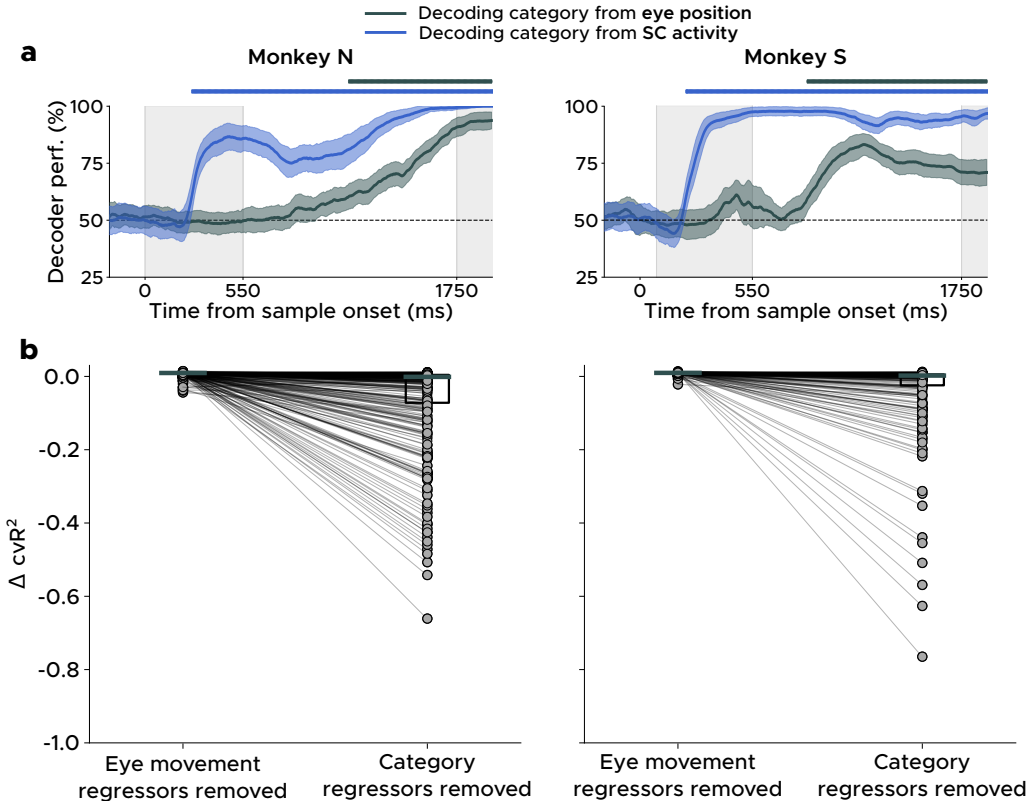

Figure 2.10: Comparison of the contribution of stimulus category and eye movements to single-neuron activity a, Time course performance of category classifiers trained on SC neural data (blue) or on horizontal and vertical eye gaze trajectories across all SC recording sessions (grey). Symbols above panels indicate time points at which classifier accuracy performed significantly above chance (P <0.05, permutation test). **b,** Change in cross-validated R^2 values from full linear encoding models (which include both eye movements and category regressors) to reduced linear models that with either eye movement-related or category regressors removed.

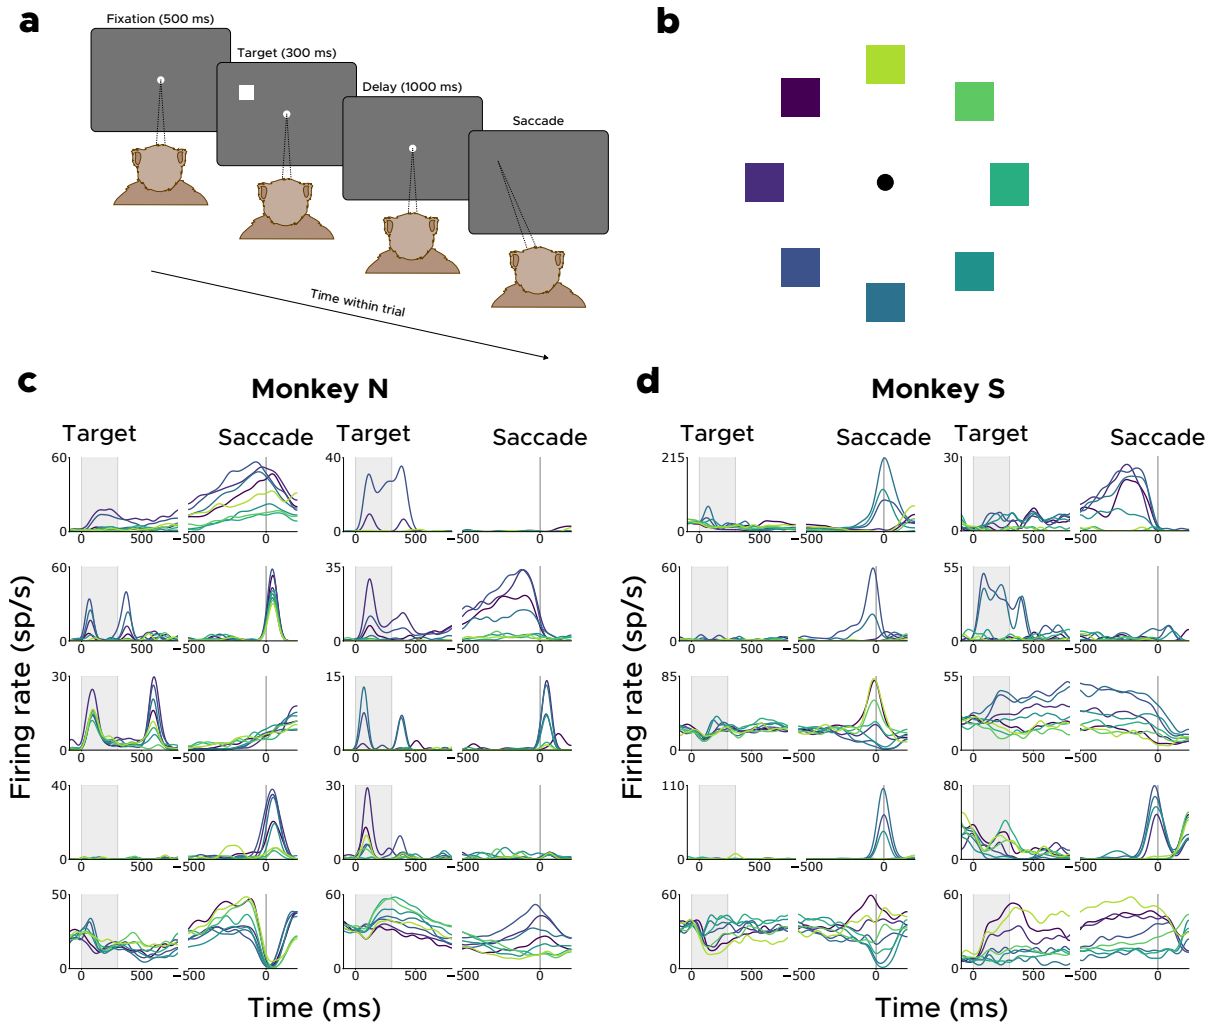

Figure 2.11: SC responses during the memory-guided saccade (MGS) task a, Schematic of the MGS task. **b,** Color-coding scheme of the eight stimulus positions used in the MGS task. **c,** Example PSTHs of SC neurons during the MGS task in Monkey N. Each trace corresponds to one of the right stimulus positions illustrated in **b**. **d,** same as *c* but for Monkey S.

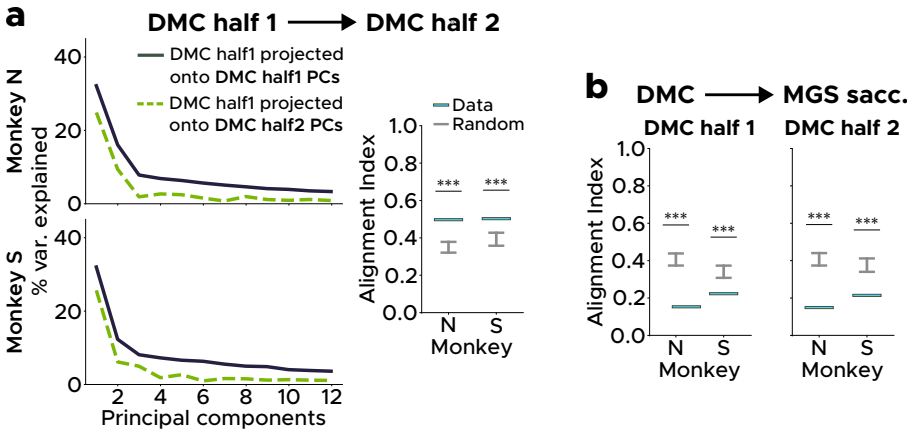

Figure 2.12: Additional subspace alignment indices a, Alignment between SC neural subspaces recorded during the first half and second half of the trials for each session. Data are significantly more aligned than chance (Monkey N: AI = .499, $P < .001$, Money S: AI = .505, $P < .001$). **b,** Left: SC neural subspaces during the first half of DMC trials and the MGS saccade period were significantly misaligned (Monkey N: AI = .155, $P < .001$, Money S: AI = .225, $P < .001$). Right: SC neural subspaces during the second half of DMC trials and the MGS saccade period were significantly misaligned (Monkey N: AI = .151, $P < .001$, Money S: AI = .216, $P < .001$).

2.5.1 Supplemental Tables

Table 2.1: SC and LIP firing rates by epoch

Monkey	Epoch	LIP firing rate (sp/s) median ± mad	SC firing rate (sp/s) median ± mad	P-value
N	Baseline	3.39 ± 2.38	2.36 ± 2.03	$.021^{*}$
N	Sample	4.45 ± 3.41	2.64 ± 2.37	$.003^{*}$
N	Early Delay	4.34 ± 2.65	2.53 ± 2.15	$<.001^{*}$
N	Late Delay	4.15 ± 2.73	2.49 ± 2.24	$.003^{*}$
N	Test	5.42 ± 3.91	3.29 ± 2.68	$<.001^{*}$
S	Baseline	1.810 ± 1.086	2.079 ± 1.573	.343
S	Sample	2.009 ± 1.246	2.368 ± 1.757	.168
S	Early Delay	1.793 ± 1.099	1.826 ± 1.369	.800
S	Late Delay	2.022 ± 1.233	1.997 ± 1.592	.912
S	Test	2.027 ± 1.353	2.397 ± 1.814	.331

* $P < .05$

** $P < .005$

* $P < .001$

Table 2.2: SC inactivation session information

Monkey	Exp	Treatment	Concentr. (µg/µL)	Injection vol. (µL)	# DMC trials (control)	# DMC trials (treatment)
N	1	Saline	8	0.50	288	292
N	2	Muscimol	5	0.33	308	222
N	3	Muscimol	5	0.33	275	353
N	4	Muscimol	5	0.33	251	510
N	5	Muscimol	5	0.25	331	456
N	6	Muscimol	5	0.25	294	385
N	7	Muscimol	5	0.25	359	218
N	8	Saline	8	0.25	273	369
S	1	Saline	8	0.25	285	427
S	2	Muscimol	5	0.25	302	278
S	3	Muscimol	5	0.25	403	104
S	4	Muscimol	5	0.25	217	138
S	5	Saline	8	0.25	258	285
S	6	Saline	8	0.25	281	547
S	7	Muscimol	5	0.25	258	164
S	8	Muscimol	5	0.25	282	182
S	9	Saline	8	0.25	305	312
S	10	Muscimol	5	0.25	290	240

Details of injection experiments in two monkeys. Each row contains an individual experimental session, and each column contains the details of that session. The column descriptions are as follows: "Monkey:" indicates which monkey the experiment was performed with; "Exp:" session number by animal; "Treatment:" indicates whether the injection was saline or muscimol; "Concentr.:" concentration of muscimol or saline (in µg/µL); "Injection vol.:" total injection volume (in µL); "N DMC trials (control):" total number of completed trials (correct and incorrect) for the DMC task during the control (pre-treatment) block; "N DMC trials (treatment):" total number of completed trials (correct and incorrect) for the DMC task during the treatment block.

Table 2.3: MGS peak saccade velocity (PSV) during SC inactivation experiments

Monkey	Exp	Contralateral PSV (°/s)			Ipsilateral PSV (°/s)		
		Ctrl ($\mu \pm \sigma$)	Inject ($\mu \pm \sigma$)	P-val	Ctrl ($\mu \pm \sigma$)	Inject ($\mu \pm \sigma$)	P-val
N	1†	434 ± 105	459 ± 96	.318	314 ± 75	300 ± 64	.370
N	2	387 ± 96	232 ± 47	$<.001^{*}$	282 ± 45	334 ± 45	$<.001^{\circledast}$
N	3	381 ± 68	237 ± 38	$<.001^{*}$	273 ± 43	329 ± 44	$<.001^{\circledast}$
N	4	447 ± 108	243 ± 51	$<.001^{*}$	260 ± 51	302 ± 46	$<.001^{\circledast}$
N	5	432 ± 94	257 ± 49	$<.001^{*}$	269 ± 54	307 ± 62	$<.001^{\circledast}$
N	6	443 ± 79	264 ± 52	$<.001^{*}$	264 ± 46	297 ± 55	$<.001^{\circledast}$
N	7	416 ± 62	270 ± 57	$<.001^{*}$	245 ± 59	310 ± 57	$<.001^{\circledast}$
N	8†	431 ± 79	421 ± 90	.579	252 ± 67	264 ± 60	.344
S	1†	273 ± 51	292 ± 65	.123	287 ± 65	299 ± 76	.317
S	2	260 ± 111	193 ± 84	$<.001^{*}$	452 ± 55	592 ± 108	$<.001^{\circledast}$
S	3	430 ± 67	174 ± 100	$<.001^{*}$	446 ± 70	494 ± 60	$<.001^{\circledast}$
S	4	457 ± 62	209 ± 109	$<.001^{*}$	464 ± 67	533 ± 89	$<.001^{\circledast}$
S	5†	444 ± 46	465 ± 30	$.008^{\circledast}$	426 ± 73	416 ± 47	.336
S	6†	435 ± 40	426 ± 43	.276	469 ± 60	458 ± 56	.280
S	7	440 ± 62	221 ± 107	$<.001^{*}$	504 ± 89	700 ± 173	$<.001^{\circledast}$
S	8	471 ± 63	185 ± 97	$<.001^{*}$	488 ± 77	571 ± 98	$<.001^{\circledast}$
S	9†	465 ± 48	465 ± 62	.982	427 ± 74	452 ± 67	.078
S	10	465 ± 64	188 ± 54	$<.001^{*}$	439 ± 75	568 ± 108	$<.001^{\circledast}$

† Saline session

* Treatment PSV < Control PSV (all $P < .050$)

⊛ Treatment PSV > Control PSV (all $P < .050$)

Peak saccade velocity (PSV) during the MGS task for the SC inactivation experiments. Rows contain individual experimental sessions and columns contain details for each session. The column descriptions are as follows: "Exp:" session number by monkey (saline injection sessions are indicated with a dagger superscript); "Contralateral PSV Ctrl:" PSV (°/s) for correct trials in which the monkey made a saccade towards one of the three targets that are in the inactivated hemifield (i.e., *contralateral* to the injected hemisphere) during the control block; "Contralateral PSV Inject:" same as "Contralateral PSV Ctrl" but for the treatment block; "P-value:" two-tailed p-value from a non-parametric permutation test (5000 permutations) to assess whether there was a significant difference in mean PSV between control and treatment blocks for the *contralateral* saccade trials; "Ipsilateral PSV Ctrl:" PSV (°/s) for correct trials in which the monkey made a saccade towards one of the three targets that are in the inactivated hemifield during the control block; "Ipsilateral PSV Inject': same as "Ipsilateral PSV Ctrl" but for the treatment block; "P-value:" two-tailed p-value from a non-parametric permutation test (5000 permutations) to assess whether there was a significant difference in mean PSV between control and treatment blocks for the *ipsilateral* saccade trials.

Table 2.4: DMC performance during SC inactivation experiments

Monkey	Exp	Control perf. ($\mu \pm \sigma$)	Inact. perf. ($\mu \pm \sigma$)	Fisher-exact odds ratio	Fisher-exact P-value
N	1†	86.81 ± 3.91	85.27 ± 4.06	1.136	.633
N	2	81.82 ± 4.31	53.15 ± 6.56	3.966	$<.001^*$
N	3	83.64 ± 4.37	50.99 ± 5.21	4.912	$<.001^*$
N	4	89.64 ± 3.77	56.08 ± 4.31	6.778	$<.001^*$
N	5	80.06 ± 4.30	67.76 ± 4.29	1.910	$<.001^*$
N	6	88.10 ± 3.70	73.25 ± 4.42	2.703	$<.001^*$
N	7	89.97 ± 3.11	45.41 ± 6.61	10.785	$<.001^*$
N	8†	92.67 ± 3.09	89.43 ± 3.14	1.495	.170
S	1†	78.95 ± 4.73	82.44 ± 3.61	0.799	.283
S	2	88.41 ± 3.61	64.03 ± 5.64	4.286	$<.001^*$
S	3	87.84 ± 3.19	70.19 ± 8.79	3.068	$<.001^*$
S	4	92.63 ± 3.48	56.52 ± 8.27	9.663	$<.001^*$
S	5†	86.05 ± 4.23	84.56 ± 4.19	1.126	.716
S	6†	93.95 ± 2.79	93.05 ± 2.13	1.159	.662
S	7	95.74 ± 2.47	75.00 ± 6.63	7.485	$<.001^*$
S	8	98.58 ± 1.38	75.82 ± 6.22	22.159	$<.001^*$
S	9†	92.46 ± 2.96	94.87 ± 2.45	0.663	.248
S	10	97.24 ± 1.89	73.33 ± 5.59	12.818	$<.001^*$

† Saline session

* Treatment perf. $<$ Control perf. (all $P < .050$)

Details of performance on the DMC task for the SC inactivation experiments. Each row contains an individual experimental session, and each column contains the details of that session. The column descriptions are as follows: "Monkey:" indicates which monkey the experiment was performed with; "Exp:" session number by animal (saline injection session are indicated with a dagger superscript); "Control perf.:" Mean $\pm$ SD accuracy (in %) for DMC trials during the control block; "Inact. perf.:" Mean $\pm$ SD accuracy (in %) for DMC trials during the treatment block; "Fisher-exact odds ratio:" odds ratio for Fisher-exact test to compare performance on control vs. treatment blocks.; "Fisher-exact P-value:" p-value for Fisher-exact test to compare performance on control vs. treatment blocks.

CHAPTER 3

CONTEXT-DEPENDENT STIMULUS ENCODING IN CORTICAL AND SUBCORTICAL VISUAL REGIONS

Barbara Peysakhovich, Stephanie M. Tetrick, Guilhem Ibos, and David J. Freedman

Author contributions: BP, GI and DJF conceived of the project. BP supervised animal training, and BP, SMT, and GI performed animal training and handling. BP and SMT collected electrophysiological recording data. BP performed all data analysis and visualization. BP and DJF wrote the manuscript.

Abstract

Neurons in the primate lateral intraparietal area (LIP) and superior colliculus (SC) reliably encode the learned categories of stimuli during visual motion categorization tasks. Moreover, neurons in the LIP, the SC, and the middle temporal are sensitive to the direction of movement of motion stimuli, even when stimuli are passively presented. However, it is unknown how task context modulates the encoding of category-related information in these areas, particularly when stimulus and category information is irrelevant for behavior. Here, we investigated how behavioral task demands affect category and sensory encoding in the SC, LIP, MT. We trained monkeys to alternate between blocks of the motion categorization task and blocks in which they passively viewed the same stimuli as for the categorization task and received a reward for maintaining fixation. The physical stimulus and stimulus location was identical in the two blocks, but only the categorization task required the monkeys to use the stimulus information to obtain a reward. We could therefore compare, in the same neurons, how behavioral context affects stimulus encoding.

We observed significantly weaker stimulus direction encoding during passive viewing than during the categorization task in all three brain areas. Moreover, although both LIP and SC encoded stimulus category during the categorization task, category encoding was largely absent during passive viewing in both areas. These results indicate that neural populations in the LIP and the SC can flexibly route sensory input based on current behavioral demands.

Keywords: middle temporal area, lateral intraparietal area, superior colliculus, cognition, categorization, direction encoding, subcortical, vision, parietal, decision making, primate, electrophysiology

3.1 Introduction

Behavioral demands and task contexts can profoundly affect how animals interact with stimuli in their environment, as well as how the brain responds to these stimuli. Previous studies have compared the neural encoding of stimuli during active task contexts in which animals are required to engage with stimuli, and contexts in animals are passively presented with the same stimuli. In primates, task engagement has been shown enhances encoding of stimulus-related information in several brain regions, including the prefrontal cortex (Hussar and Pasternak, 2009, 2012, 2010), V4 (Popovkina and Pasupathy, 2022; Zamarashkina et al., 2020), and middle temporal area (MT) (Scott et al., 2022). In ferrets trained on a tone detection task, transition from an active detection task to a passive listening task rapidly changes the spectrotemporal response properties of neurons in the primary auditory cortex (Fritz et al., 2003; Fritz, 2005) and inferior colliculus (Shaheen et al., 2021; Slee and David, 2015). Task engagement is closely related to attention, which has been shown to have widespread effects on neural responses, including enhanced neural responses, increased discriminability between stimuli, and reduced variability (Martinez-Trujillo and Treue, 2002;

Treue and Maunsell, 1996; Martinez-Trujillo and Treue, 2004; Cohen and Newsome, 2008; Reynolds and Chelazzi, 2004; Yantis and Serences, 2003; Maunsell and Cook, 2002; Cohen and Maunsell, 2009).

In this study, we investigated how behavioral demands affect encoding of sensory and cognitive information in the lateral intraparietal area (LIP), superior colliculus (SC), and middle temporal area (MT). The LIP, located located in the lateral bank of the intraparietal sulcus (Blatt et al., 1990), has been shown to exhibit binary-like category tuning during visual categorization tasks (Swaminathan et al., 2013; Swaminathan and Freedman, 2012; Sarma et al., 2016; Mohan et al., 2021; Freedman and Assad, 2006; Rishel et al., 2013). The SC is a midbrain structure that is involved in spatial orienting and has been recently shown to play a causal role in abstract visual categorization (Peysakhovich et al., 2023) (as discussed in *Chapter 2*). The middle temporal area (MT) is an the extrastriate visual region located on the posterior bank of the superior temporal sulcus (Maunsell and van Essen, 1983a). MT receives direct input from the primary visual cortex (V1) (Cragg, 1969), and a large percentage of neurons in MT respond selectively to the direction of motion stimuli ((Maunsell and Van Essen, 1983b)).

We compared the activity of the same populations of neurons in MT, LIP, and SC across two tasks (a motion categorization task and a passive viewing task) in which animals viewed the same physical stimulus, but only one of the tasks required the monkeys to use the stimulus information to obtain a reward. We trained two monkeys to perform a delayed match-to-category (DMC) task in which they grouped 360° of motion directions into two categories based on an arbitrary category rule. The categories were defined by two perpendicular boundaries that produced four 90°-wide quadrants (**Fig. 3.1a**). To disambiguate neuronal encoding of direction vs. category, opposite quadrants were assigned to the same category, so that motion directions that are 180° apart belonged to the same category while nearby directions were often in different categories. Monkeys viewed sample and test dot-motion

stimuli separated by a 1.2 sec memory delay (**Fig. 3.1b**). On each trial, they received a fluid reward for releasing a manual touch bar when the category of the test stimulus matched the category of the sample stimulus (Match trials). If the test stimulus category did not match the sample stimulus category (Non-match trials), the monkeys were shown a second test stimulus, which always matched the sample category (and required release of the touch bar in order to receive a reward).

Monkeys also performed interleaved blocks of a passive viewing (PV) task during the same recording sessions as the DMC task. In the PV task, the monkeys viewed 3-5 motion stimuli presented for 400 ms each and separated by a 200-ms delay (**Fig. 3.1c**). Monkeys received a visual cue (a blue fixation circle instead of a white fixation circle as for the DMC task) to indicate the start of a PV trial. For both tasks, the monkeys were required to maintain gaze on a central fixation spot throughout the trial (see *Methods*). For the PV task, the monkeys received a reward for maintaining fixation on the central cue for the duration of the trial. The stimuli used in the PV task were identical to those shown in the DMC task and were presented in the same peripheral location as for the DMC task. In the DMC task, the monkeys needs to use the information provided by the stimulus (i.e., the direction and category of the stimulus) to guide behavior and make decisions, while in the PV task, the direction and category of the stimulus is irrelevant for obtaining a reward.

Following long-term training, both monkeys performed the DMC task with >85% mean accuracy during neural recording sessions (Monkey N: 89.76 ± 3.29%, Monkey S: 88.14 ± 3.57%; **Fig. 3.1c**). Monkeys performed similarly on Match trials in which sample and test stimuli were in the same vs. opposite quadrants (Monkey N: Same quad. = 88.30 ± 4.21%, Opp. quad. = 87.06 ± 4.85%, $P = .124$, Monkey S: Same quad. = 85.08 ± 6.85%, Opp. quad. = 84.66 ± 5.29%, $P = .778$, permutation test; **Fig. 3.1d**), indicating that they learned to categorize stimuli across opposite quadrants.

We recorded spiking activity during the DMC and PV tasks from 434 MT neurons (Mon-

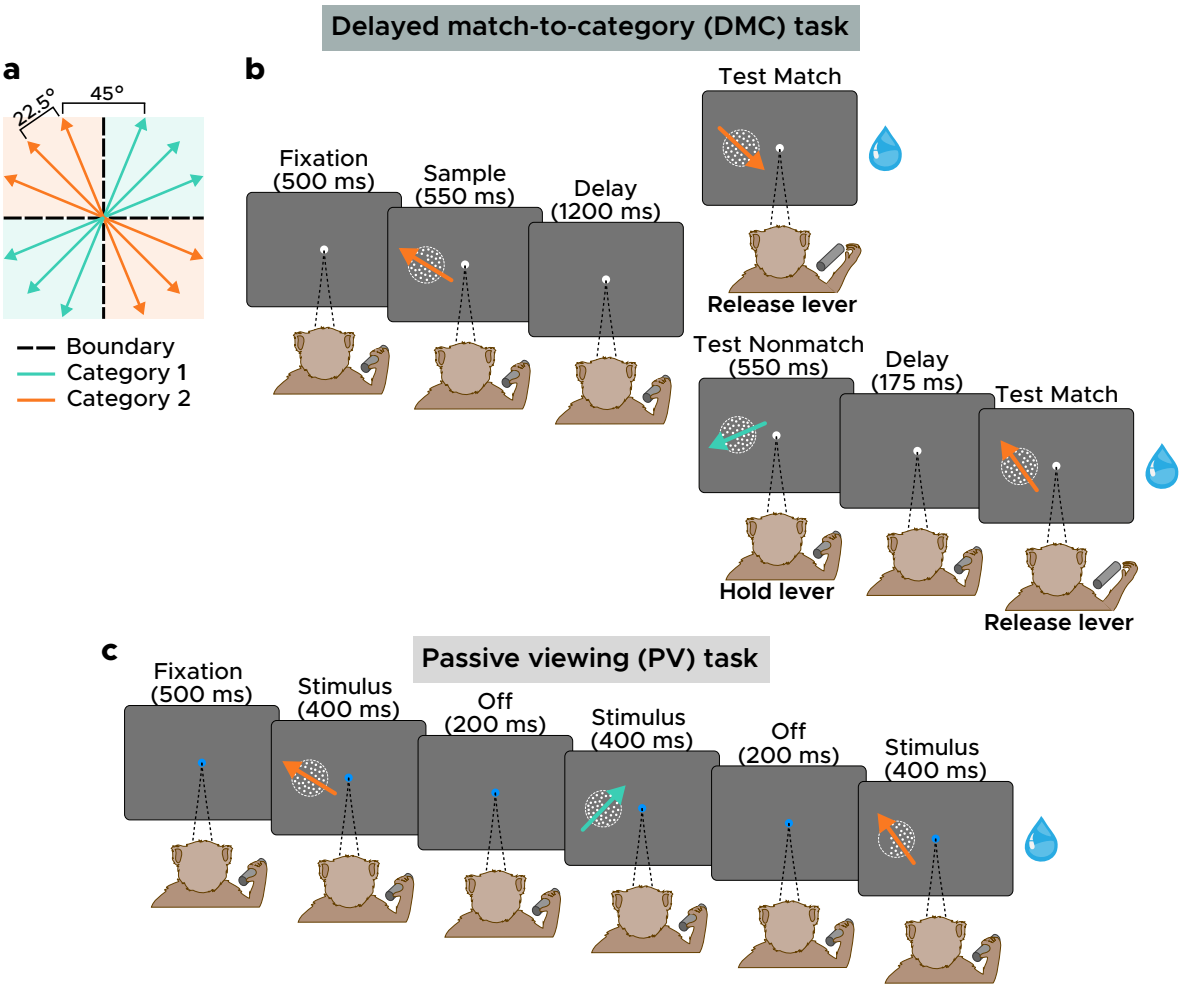

Figure 3.1: Schematics of the DMC and PV paradigms. **a,** Stimulus geometry of the two-boundary delayed match-to-category (DMC) task. 12 directions of motion are grouped into two categories based on two orthogonal category boundaries (dashed lines), such that directions that are 180° apart belong to the same category. Directions within the same quadrant are 22.5° apart, and near-boundary directions are 22.5° from the boundary. **b,** Trial structure of the DMC task. Monkeys were required to maintain gaze within a small window centered on a central fixation cue, and reported their decisions with a hand movement (holding or releasing a lever). **c,** Trial structure of the passive viewing (PV) paradigm. Monkeys view a series of 3-5 motion stimuli presented for 400 ms each and separated by 200 ms.

key N: 215, Monkey S: 219), 496 LIP neurons (Monkey N: 169, Monkey S: 327) and 539 SC neurons (Monkey N: 312, Monkey S: 227). The neurons recorded in LIP and SC are a subsample of neurons described in *Primate superior colliculus is engaged in abstract higher-order cognition*. Because recorded data from all neurons during both tasks, we could compare

response properties between the two tasks on a neuron-by-neuron basis.

3.2 Results

3.2.1 Differences in firing rates between the DMC and PV task

We first compared the firing rates of individual neurons during the DMC and PV task. For each neuron, we calculated the mean firing rate for both tasks for each stimulus direction, as well as across all directions. We then computed the difference in firing rate between tasks (FR_{DMC}-FR_{PV}) and determined whether the resulting distribution of difference values across neurons was shifted significantly above or below 0 (with shifts above 0 indicating higher firing rates during DMC task). In all three brain areas, the distribution of FR_{DMC}-FR_{PV} values was significantly above 0 when considering both the mean responses across all directions (median ± m.a.d: MT = .061 ± .273, LIP = .071 ± .337, SC = .061 ± .283; all $P < .001$, permutation test; **Fig. 3.2 a-b**) and the responses for each neuron's preferred direction (identified during the DMC task; (median ± m.a.d: MT = .216 ± .373, LIP = .253 ± .435, SC = .315 ± .416; all $P < .001$, permutation test.; **Fig. 3.2 c-d**). However, in all brain areas, the firing rate difference distribution were significantly below 0 for the least-preferred direction (median ± m.a.d: MT = -.124 ± .399, LIP = -.169 ± .502, SC = -.184 ± .551; all $P < .001$, permutation test; **Fig. 3.2 e-f**), indicating higher firing rates during the PV task than the DMC task. In **Fig. 3.2g**, we plotted the firing rate modulation index ([DMC-PV]/[DMC+PV]; positive values indicate higher DMC firing rates) as a function of preferred direction rank, revealing that responses during the PV task are most attenuated for preferred directions, while responses for least-preferred directions are amplified.

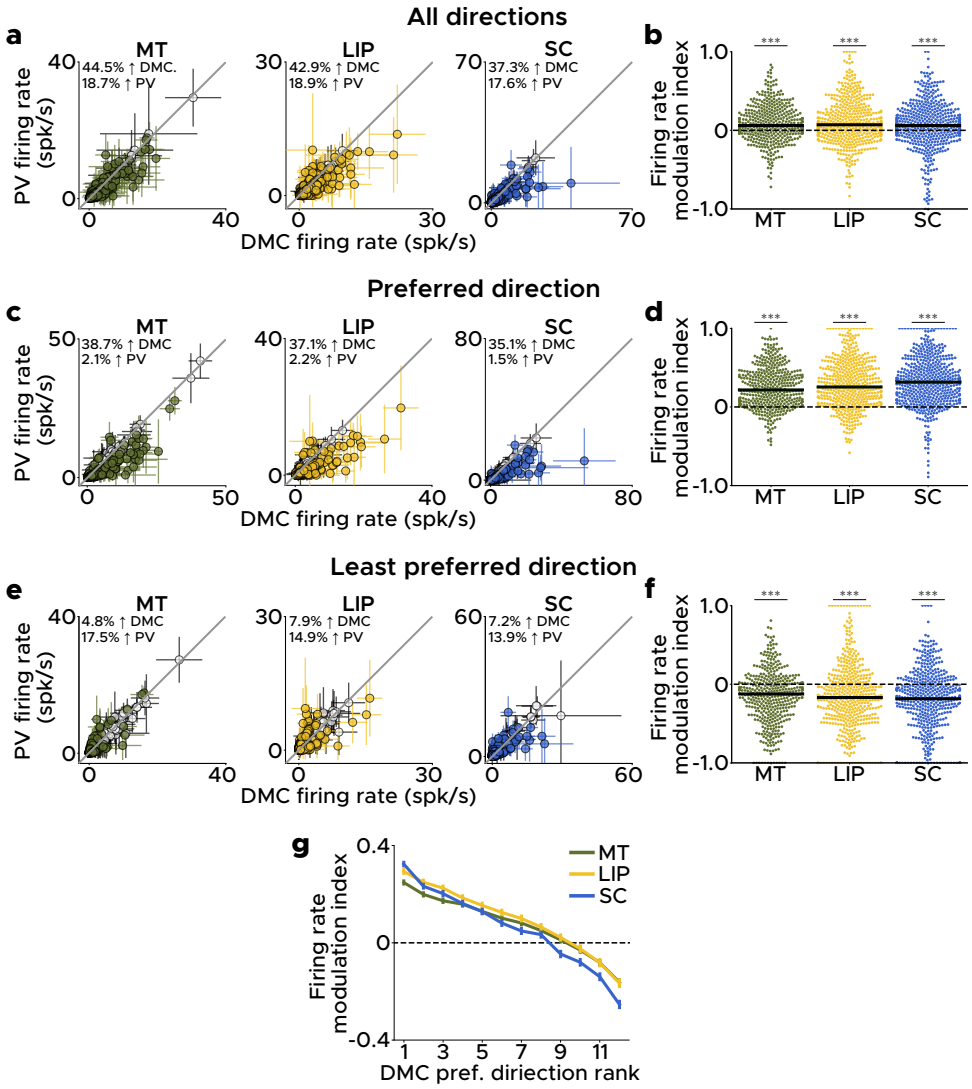

Figure 3.2: Task context modulates firing rates in MT, LIP, and SC neurons. a, Scatter plots showing mean firing rates during the DMC (x-axis) vs. PV (y-axis) task for all neurons in MT (left), LIP (middle), and SC (left). Error bars indicate s.d. across trials per neuron. Text in the top left corner of each subplot indicates the % of cells that have significantly higher firing rates during the DMC task than the PV task, and the % of cells that have significantly higher firing rates during the PV task. **b,** Distribution of firing rate modulation index values (DMC-PV/DMC+PV) across neurons in each area. Horizontal solid black lines indicate median. **c, d** Same as *a* and *b* but for each neuron's preferred direction (defined during the DMC task). **e, f** Same as *a* and *b* but for each neuron's least-preferred direction (defined during the DMC task). **g,** Mean firing rate modulation index as a function of preferred direction rank (defined during the DMC task). *** indicates $P < .001$.

3.2.2 Neurons in MT, LIP and SC encode stimulus direction more reliably during the categorization task than the passive viewing task

We then investigated the direction tuning properties of neurons during the DMC and PV task in all three brain areas by comparing the peristimulus time histograms (PSTHs) and mean response tuning curves in individual neurons. Many MT neurons showed strong motion direction tuning, with high response rates for a preferred direction and firing rates decreasing as a function of distance from the preferred direction (**Extended Fig. 3.10**). Some MT neurons had remarkably similar response tuning curves in during the DMC and PV tasks (**Extended Fig. 3.10 a-b**), although many MT neurons showed considerably weaker direction tuning during the PV task compared to during the DMC task; some neurons had markedly lower firing rates for their preferred direction during the PV vs. DMC task (**Extended Fig. 3.10 c-d**), and others had both lower firing rates for their preferred direction and increased firing rates for non-preferred directions (**Extended Fig. 3.10 e-f**). Some LIP and SC neurons were also direction-tuned during the DMC task, and many of these neurons showed markedly reduced (or absent) direction tuning during the PV task (**Extended Fig. 3.11** and **Extended Fig. 3.12**).

Single-neuron direction encoding during the DMC and PV tasks

To quantify the strength of direction encoding in each neuron during the visual epoch (from 50 to 400 ms after stimulus onset), we used three complementary methods: (1) a receiver operating characteristic (ROC)-based method in which we quantified how discriminable each stimulus direction is based on a neuron's responses by computing area-under-the-ROC-curve values for each of the 12 stimulus direction vs. all other directions and calculating the mean and maximum ROC across conditions for each neuron; (2) the mutual information (MI) between spike counts and motion direction, which quantifies how much information (in bits) a neuron's activity provides about stimulus direction; and (3) a direction tuning index

(DTI) that quantifies (on a scale of 0 to 1) the magnitude of a neuron's unimodal direction tuning based on the circular variance method introduced in (Mazurek et al., 2014). Neurons that have unimodal, cosine-like response tuning curves have DTI close to 1, while neurons that have either uniform tuning curves or bimodal tuning curves with similar responses to directions that are 180° apart (i.e., category-tuned neurons during the DMC task) have DTI close to 0. The MI and ROC methods are applied to distributions of trial-by-trial firing rates and are thus sensitive to differences in trial-to-trial variability, while the DTI method is applied on trial-average responses for each motion direction. See *Methods* for additional details about the three direction tuning measures. We identified neurons that had significantly elevated DTI, direction ROC, and MI values (relative to null distributions of shuffled values for each neuron obtained by permuting direction labels for each trial; see *Methods*). **Extended Table 3.1** shows the percentage of neurons in each brain area with significantly above-chance DTI, direction ROC, and MI values, as well as the percentage of neurons with significantly higher direction-encoding values in the DMC vs. PV task.

To investigate whether there are differences in the strength of direction encoding during the DMC vs. PV task in each brain area, we computed the difference in direction ROC and MI between tasks for each neuron (e.g., $\mathrm{ROC_{DMC}}$-$\mathrm{ROC_{PV}}$) and determined whether the resulting distribution of difference values across neurons was shifted significantly above or below 0 (with shifts above 0 indicating stronger direction encoding during the DMC task). In all three brain areas, the distribution of $\mathrm{ROC_{DMC}}$-$\mathrm{ROC_{PV}}$ values was significantly above 0 for both the mean ROC across stimulus directions for each neuron (median ± m.a.d: MT = .006 ± .018, LIP = .007 ± .020, SC = .005 ± .017; all $P < .001$, permutation test; **Fig. 3.3 a-b**) and the maximum ROC across directions (median ± m.a.d: MT = .013 ± .046, LIP = .014 ± .051, SC = .014 ± .045; all $P < .001$, permutation test; **Fig. 3.3 c-d**). In all three brain areas, the distribution of $\mathrm{MI_{DMC}}$-$\mathrm{MI_{PV}}$ values was also significantly above 0 (median $\mathrm{MI_{DMC}}$-$\mathrm{MI_{PV}}$ ± m.a.d: MT = .045 ± .155, LIP = .010 ± .116, SC = .034 ± .116; MT and

SC: $P < .001$, LIP: $P = .013$, permutation test; **Fig. 3.3 e-f**). These results indicate that in a majority of neurons in all three brain areas, direction encoding is stronger during the DMC task than the PV task.

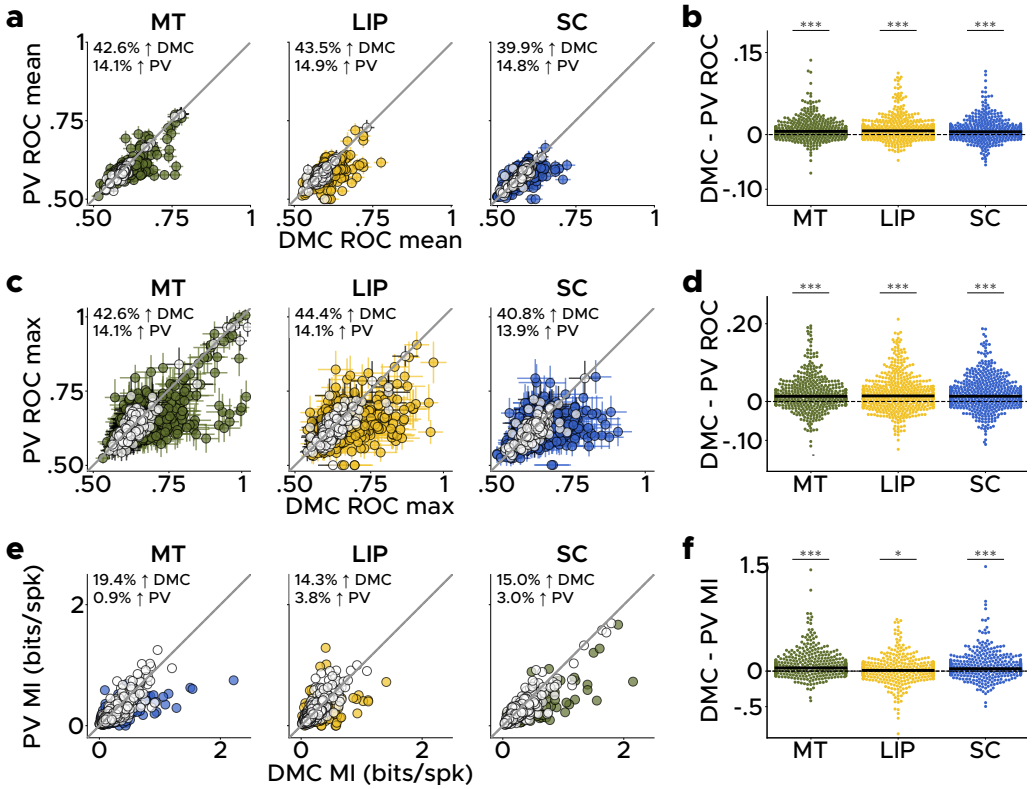

Figure 3.3: Stronger single-neuron direction encoding during the DMC task than the PV task. **a,** Scatter plots showing mean ROC (across stimulus directions) for each neuron during the DMC task (x-axis) vs. the PV task (y-axis) in MT (left), LIP (middle), and SC (left). Filled circles: neurons in which values differ significantly between tasks. Unfilled circle: neurons in which values do not differ significantly between tasks. Error bars indicate s.d. across trials for each neuron, computed via bootstrapping. Text in the top left corner of each subplot indicates the % of cells that have significantly higher ROC during the DMC task than the PV task and the % of cells that have significantly higher ROC during the PV task. **b,** Distribution of ROC differences values (DMC-PV) across neurons in each area. Horizontal solid black lines indicate median. **c, d,** Same as *a* and *b* but for maximum ROC across stimulus directions for each neuron. **e, f** Same as *a* and *b* but for mutual information between spike count and motion direction. * indicates $P < .05$, *** indicates $P < .001$.

To investigate whether there are differences in the strength of unimodal/cosine-like direction tuning during the DMC and PV tasks, we computed the difference in DTI between tasks for each neuron. In all three brain areas, the distribution of DTI_{DMC}-DTI_{PV} values was

significantly above 0 in direction-tuned neurons (i.e., neurons that had significantly above-chance DTI during the DMC task; median ± m.a.d: MT = .052 ± .097, LIP = .040 ± .097, SC = .049 ± .127; $P < .001$ in MT, LIP and SC, permutation test; **Fig. 3.4 a-b**). However, in all three brain areas, the distribution of DTI_{DMC}-DTI_{PV} values was significantly below 0 in direction un-tuned neurons, indicating weaker direction tuning during the DMC task in these neurons (MT = -.012 ± .071, LIP = -.013 ± .085, SC = -.007 ± 0.083; all $P < .001$, permutation test; **Fig. 3.4 c-d**). Neurons can have low DTI (even if ROC/MI values are high) if their response tuning curves are bi- or multimodal. The direction-untuned neurons may therefore have lower DTI during the DMC task if their tuning curves become more bimodal due to an increase in category tuning strength during the DMC task.

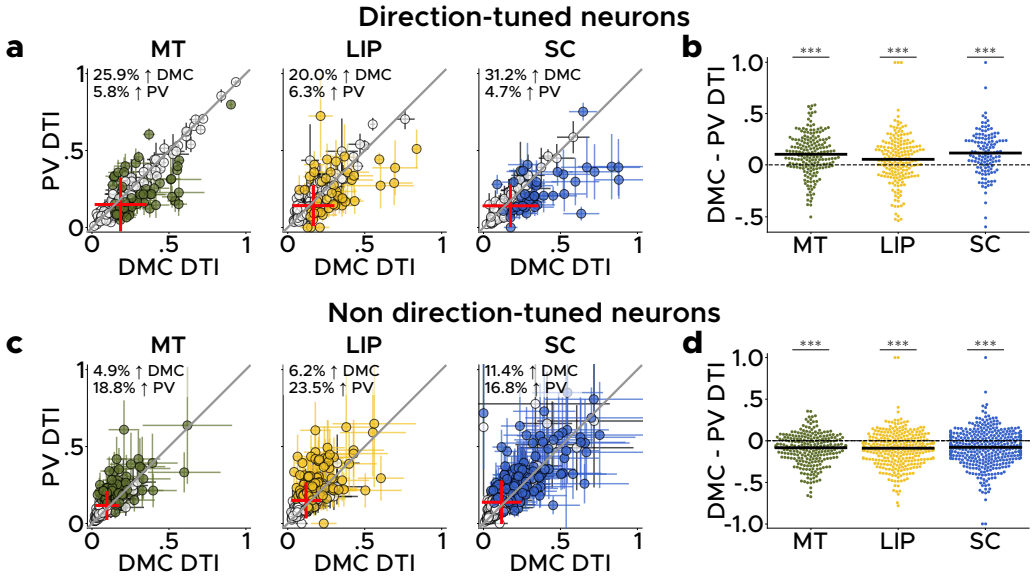

Figure 3.4: Changes in unimodal direction tuning strength between the DMC and PV task. a, Scatter plots showing mean direction tuning indices (DTI) during the DMC task (x-axis) vs. the PV task (y-axis) for all significantly direction-tuned neurons in MT (left), LIP (middle), and SC (left). Filled circles: neurons in which values differ significantly between tasks. Unfilled circle: neurons in which values do not differ significantly between tasks. Error bars indicate s.d. across trials for each neuron, computed via bootstrapping. Text in the top left corner of each subplot indicates the % of cells that have significantly higher DTI during the DMC task than the PV task and the % of cells that have significantly higher DTI during the PV task. **b,** Distribution of DTI difference values (DMC-PV]) across neurons in each area. Horizontal solid black lines indicate median. **c, d,** Same as *a* and *b* but for neurons untuned neurons. *** indicates $P < .001$.

subsubsectionPopulation-level direction encoding during the DMC and PV tasks

To quantify population-level encoding of direction in MT, LIP and SC, we trained linear support vector machine (SVM) classifiers to decode sample direction from neural population activity. The classifiers were trained on a subset of trials per direction for each neuron and tested on held-out trials (see *Methods*). We applied the direction classifiers to data across time in the trial (**Fig. 3.5a**) as well as to epoch-average spike rates (from 50- to 400 ms after stimulus onset) for both tasks (**Fig. 3.5b**). In all three brain areas, direction classifiers trained on DMC data significantly outperformed classifiers trained on PV data (MT: DMC = 92.3 (μ) ± 7.7% (s.d.), PV = 70.0 ± 12.3%; LIP: DMC = 69.6 ± 12.3%, PV = 32.8 ± 13.3%; SC: DMC = 25.5 ± 12.5%, PV = 19.0 ± 10.9, all $P < .001$, permutation test). Performance was significantly higher for classifiers trained on DMC vs. PV data even when we included only neurons that had significant DTI direction tuning during either task (MT: $N = 189$ neurons, DMC = 94.3 ± 7.0%, PV = 73.9 ± 12.3%, $P = .002$; LIP: $N = 174$, DMC = 71.1 ± 12.4%, PV = 35.7 ± 14.0%, $P < .001$; SC: $N = 144$, DMC = 30.3 ± 14.2%, N = X neurons, PV = 20.9 ± 11.58, $P = .014$, permutation test; **Extended Fig. 3.14**).

Correlations between DMC-PV tuning curves are strongest in area MT

We also characterized how correlated each neuron's tuning curves are between the DMC and PV tasks. We applied the Spearman rank correlation to the two vectors containing a neuron's mean firing rate for each of the twelve motion stimuli, separately for direction-tuned and untuned neurons (based on DTI). DMC and PV tuning curves were significantly more correlated in direction-tuned MT neurons than in tuned LIP and SC neurons (MT vs. LIP: $P < .001$; MT vs. SC: $P < .001$, permutation test), a tuning curves were more correlated in tuned LIP neurons than tuned SC neurons. ($P = 0.013$, permutation test). In all three brain areas, correlations between DMC/PV tuning curves were significantly stronger in direction-tuned vs. untuned neurons (median r ± m.a.d.: MT tuned = .511 ± .389, MT untuned =

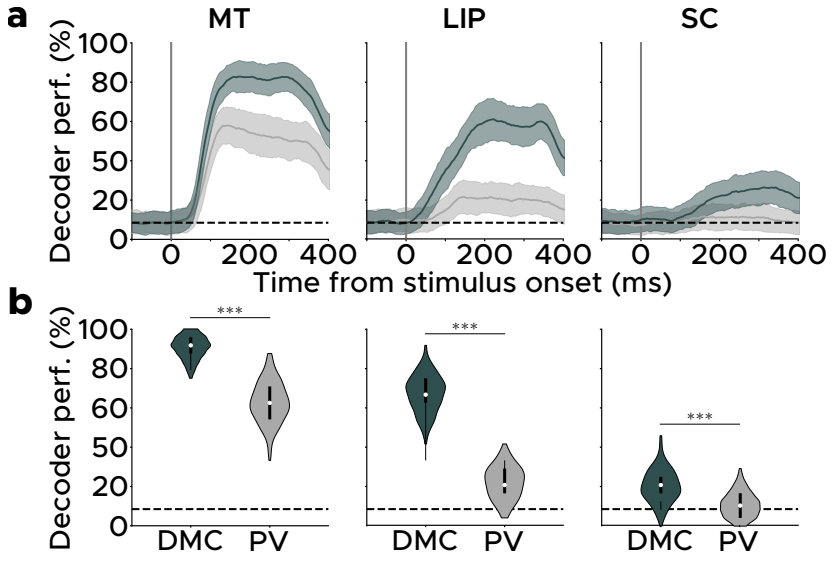

Figure 3.5: Stronger population-level direction encoding during the DMC task than the PV task. **a,** Time course of mean accuracy of stimulus direction classifiers trained on DMC (dark grey) or PV (light grey) trials for MT (left), LIP (middle), and SC (right) neural populations. Shading indicates s.d. across bootstraps. **b,** Accuracy of direction classifiers trained on whole-epoch data (from 50 to 400 ms after stimulus onset). *** indicates $P < .001$. In *e* and *f*, horizontal dashed lines indicate chance level.

.084 ± .332; LIP tuned = .362 ± .381 LIP untuned = .112 ± .342; SC tuned = .224 ± .402 SC untuned = .036 ± .337; all $P < .001$, permutation test; **Fig. 3.6**). We observed similar results when we identified neurons as direction-tuned based on the ROC direction tuning measure (**Fig. 3.15**).

Taken together, these results reveal that task demands deferentially affect stimulus direction encoding in MT, LIP, and SC. During the DMC task (when direction information is necessary for obtaining a reward), neuronal responses carry more information about stimulus direction and direction-untuned neurons show weaker cosine-like direction tuning compared to the passive context in which stimulus direction is irrelevant for behavior. Correlations between DMC and PV tuning curves were strongest in MT, followed by LIP, and were stronger in direction-tuned than untuned populations of neurons in each area.

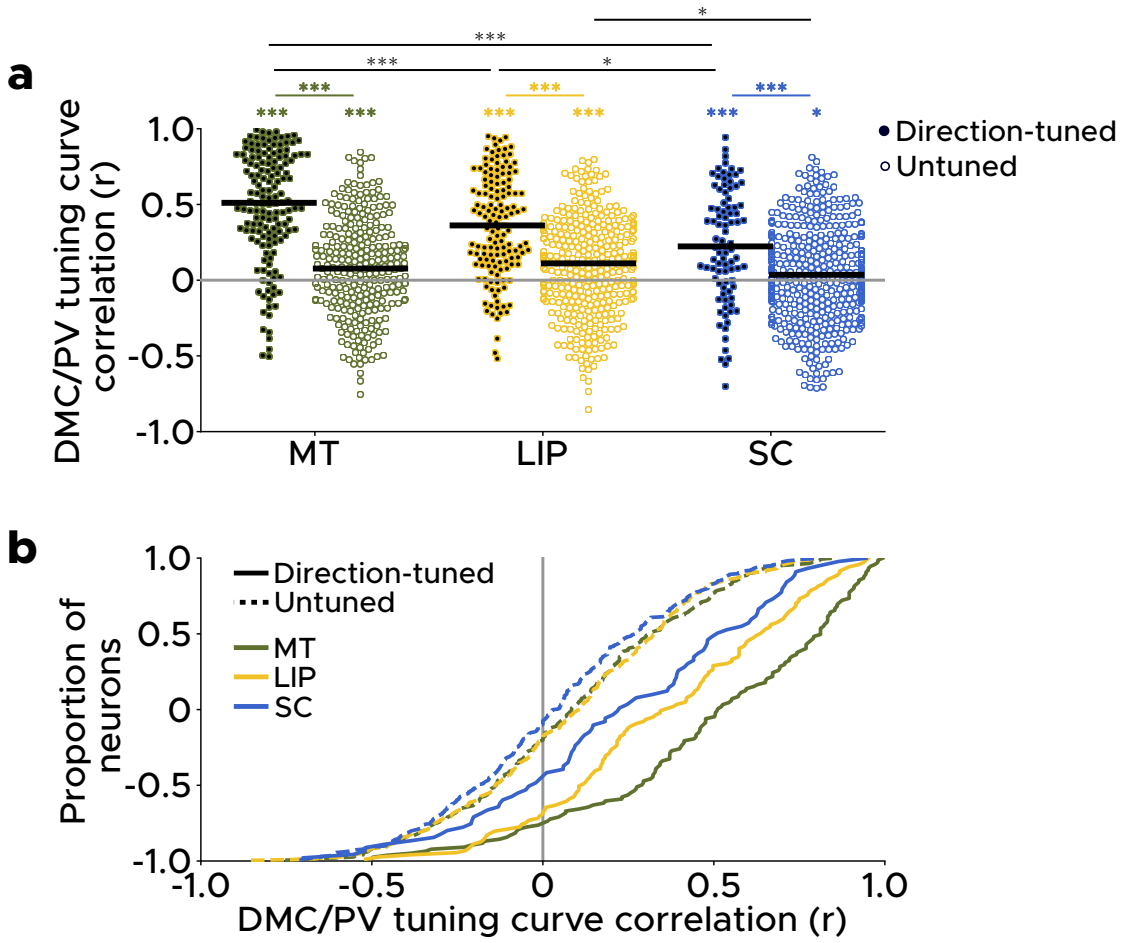

Figure 3.6: Stronger correlations between DMC-PV tuning curves in MT than in LIP and SC. a, Distributions of tuning curve correlation r values for neurons that are significantly direction-tuned (filled circles) and untuned to direction (unfilled circles) in each brain area. Horizontal dashed line indicates an r value of 0. Horizontal solid black lines indicate the median for each distribution. **b,** Empirical cumulative distribution plots for direction-tuned neurons (solid traces) and untuned neurons (dashed traces) in each brain area. Vertical gray line indicates an r value of 0. * indicates $P < .05$, *** indicates $P < .001$.

3.2.3 Category encoding is largely absent in LIP and SC during the passive viewing task

Single-neuron category tuning during the DMC and PV tasks

Next, we investigated differences in stimulus category encoding during the DMC and PV tasks. As reported in *Primate superior colliculus is engaged in abstract higher-order cogni-*

tion, many neurons in LIP and SC exhibited strong, binary-like category tuning during the DMC task, with distinct activity for directions in different categories and similar activity for directions in the same category (**Extended Fig. 3.16** and **Extended Fig. 3.17**). During the PV task, many of these neurons were not selective for category and had uniform tuning curves; in some neurons, mean firing rates during the PV task were similar to firing rates for the non-preferred category during the DMC task (**Extended Fig. 3.11c-f**, **Extended Fig. 3.12c-d**), and in others, mean firing rates during the PV task were similar to firing rates for the preferred category during the DMC task (**Extended Fig. 3.11f**, **Extended Fig. 3.12e-f**).

We quantified the strength of category tuning in individual neurons during both tasks using an ROC-based category tuning index (rCTI) that compares neuronal discrimination between pairs of directions in the same vs. different categories (see *Methods*). rCTI can range from -0.5 to 0.5, with positive values indicating larger differences in firing rates between pairs of directions in different vs. same categories (and thus strong category tuning) and negative values indicating the opposite. We computed the rCTI for each neuron using the mean response across the late visual epoch (from 200 to 400 ms after stimulus) for each trial. We then compared rCTI values between the DMC and PV task. In MT, there was no difference in rCTI between tasks (mean $\mathrm{rCTI_{DMC}}$-$\mathrm{rCTI_{PV}}$ = .000, s.d. = .021, P = .675, permutation test), as expected given that category tuning is largely absent in MT during the visual epoch of the DMC task. In LIP and SC, the distributions of $\mathrm{rCTI_{DMC}}$-$\mathrm{rCTI_{PV}}$ values across neurons were significantly above 0 (LIP = .005 ± .026; SC = .006 ± .033; both P < .001, permutation test; **Fig. 3.7a**, indicate stronger single-neuron category tuning during the DMC task.

Single-neuron quadrant tuning during the DMC and PV tasks

However, only a small proportion of neurons in LIP and SC showed significantly elevated epoch-averaged rCTI even during the DMC task, indicating that rCTI may be too strict of a measure for detecting subtle changes between tasks in a neuron's category selectivity. We therefore developed an additional measure to quantify single-neuron quadrant tuning, or selectivity for the three directions in one of the four quadrants of motion directions. For each of the eight near-boundary directions, we compared the similarity of neural responses for that reference direction and for the two directions that are 45° away, one of which is in the same quadrant (and therefore the same category as the reference direction), and the other of which is in a different quadrant **Fig. 3.7b**. We computed the difference in ROC between the same-quadrant and different-quadrant pairs (ROC_{diff}-ROC_{same}, as well as the differences in firing rates between the two pairs (FR_{diff}-FR_{same}). We computed a quad-ROC and quad-FR value for each neurons by average the ROC and firing rate difference values across the eight near-boundary conditions. A quadrant-tuned neuron will consistently respond more similarly to same-quad direction pairs (i.e., ROC_{same} will be lower than ROC_{diff}) and will therefore have above-zero quad-ROC and quad-FR values.

To quantify differences in single-neuron quadrant tuning between the DMC and PV tasks, we computed the difference in quad-ROC and quad-FR between tasks for each neuron. In LIP and SC, the distributions of quad-ROC_{DMC}-quad-ROC_{PV} and quad-FR_{DMC}-quad-FR_{PV} values were significantly above 0 (LIP: ROC = .004 ± .016, $P < .001$, FR = .416 ± .975, $P < .001$; SC: ROC = .010 ± .001, $P = .013$, FR = .364 ± .973, $P = .002$, permutation test; **Fig. 3.7 c-d**), indicating stronger quadrant during the DMC task than the PV task. In MT, the distribution of quad-ROC differences values was not significantly different from 0 (median = .000 ± .016, $P = .535$, permutation test), although the distribution of quad-FR rates was significantly above 0 (median = .285 ± .877, $P = .003$, permutation test).

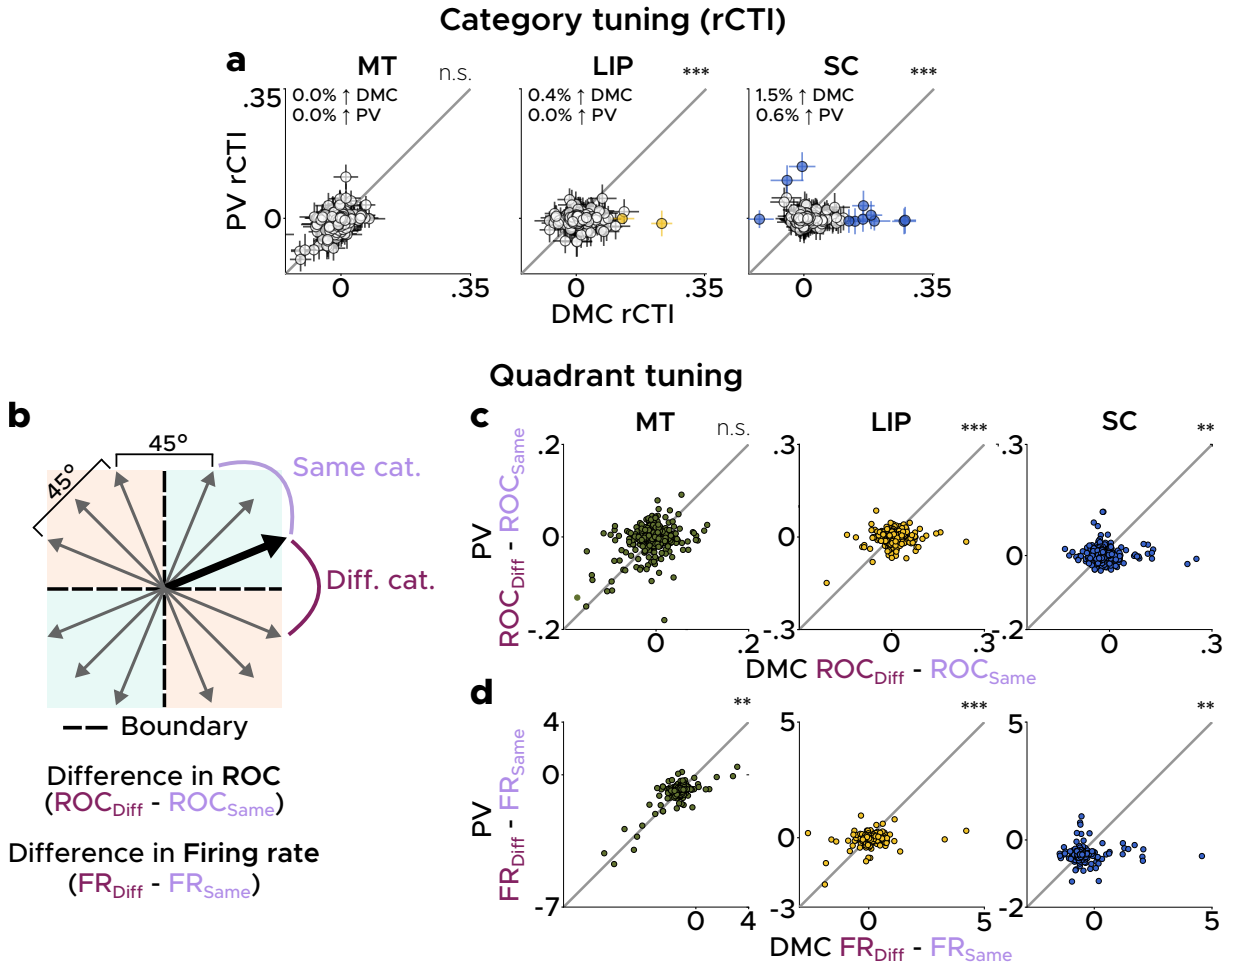

Figure 3.7: Comparison of single-neuron category tuning during the DMC and PV tasks. a, Scatter plots showing mean category tuning indices (rCTI) during the DMC task (x-axis) vs. the PV task (y-axis). **b,** Schematic of the quadrant-tuning analysis. We compute the ROC and difference in firing rate between the example direction in bold and the two highlighted directions, and substract the ROC and FR difference for the same-quadrant pair from the different-quadrant pair. To compute a single quad-ROC and quad-FR value per neuron, we average the values across all eight near-boundary motion directions **c,** Scatter plots showing mean wuad-ROC during the DMC task (x-axis) vs. the PV task (y-axis). **d,** same as *c* but for quad-FR values.

Population-level category encoding during the DMC and PV tasks

We next used linear support vector machine (SVM) classifiers to quantify the amount of category encoding in the MT, LIP and SC neural populations during the DMC and PV tasks. To evaluate the strength of category encoding in a direction-independent manner, the

classifiers were trained on trials from two quadrants (one from each category) and validated on the remaining two quadrants (see *Methods*), as the will be able to generalize between the two quadrants within the same category if neural populations encode category in a binary-like manner. In addition, this approach prevents direction tuning from contributing to category decoding performance by decorrelating direction and category between the sample and test sets. Note that we find below-chance classifier performance when the population shows strong direction tuning.

In MT, there was no difference in performance between cross-quadrant category classifers trained on DMC or PV data; for the two tasks, classifer performance was similarly below below chance, indicating that direction signals dominated population activity (DMC = 17.01 ± 8.27%, PV = 15.57 ± 7.11%; $P = .327$, permutation test). In LIP and SC, DMC classifier performance was significantly higher than that of the PV classifiers (LIP: DMC = 62.10 ± 10.35%, PV = 37.11 ± 5.54%; SC: DMC = 92.96 ± 3.99%, PV = 56.38 ± 6.71%, both $P < .001$, permutation test; **Fig. 3.8b**). We performed the same category decoding analysis on a subpopulation of neurons that excluded all significantly direction-tuned neurons (based on DTI) to minimize the effect of direction encoding on classifier performance (**Fig. 3.8c**). In MT, accuracy of category classifiers trained only non-tuned neurons was at chance level and not significantly different for DMC vs. PV (DMC = 53.38 ± 6.33%, PV = 51.16 ± 6.57%; $P = .549$, permutation test). IN LIP and SC, results were similar to those for classifiers trained on all neurons; performance was significantly higher for classifiers trained on DMC vs. PV data (LIP: DMC = 70.51 ± 5.41%, PV = 54.63 ± 6.71%; SC: DMC = 79.57 ± 5.35%, PV = 60.51 ± 5.40%), both $P < .001$, permutation test).

3.3 Discussion

In this study, we investigated how task demands affect responses of neurons in MT, LIP, and SC. We trained monkeys to alternate between blocks of two tasks: (1) a visual cate-

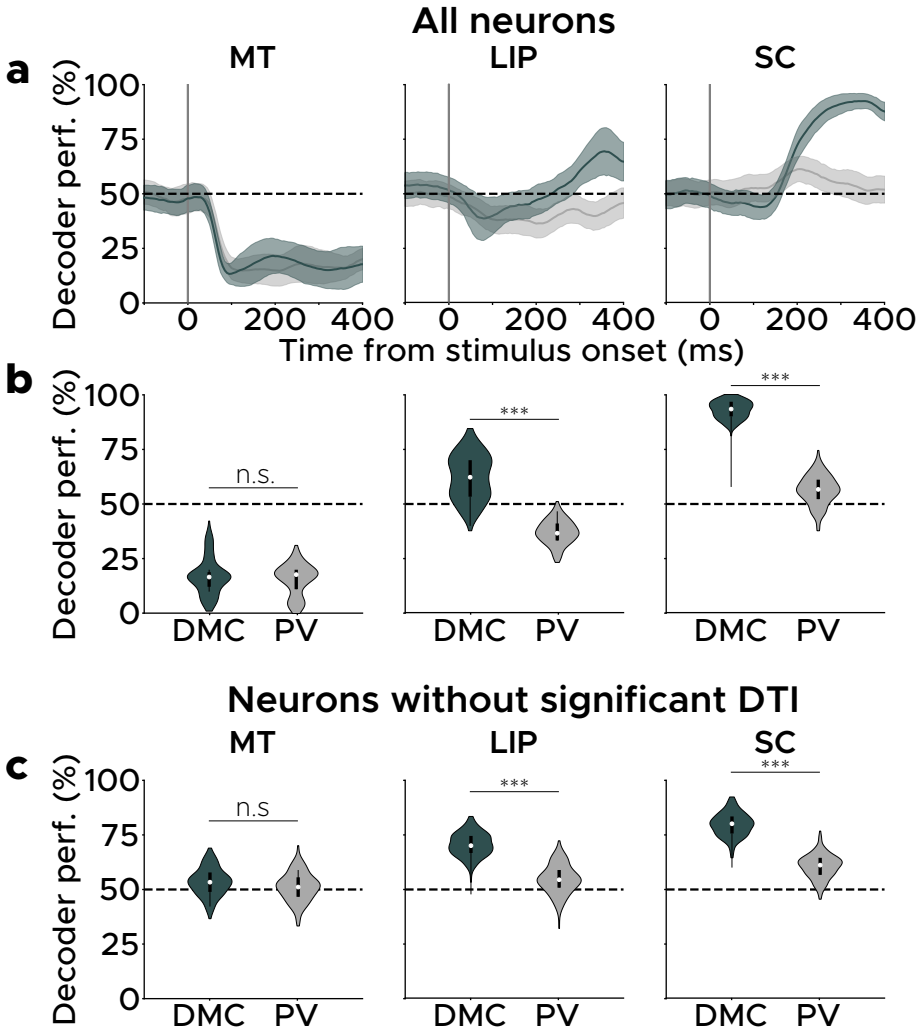

Figure 3.8: Stronger population-level category encoding during the DMC task compared to the PV task. a, Time course of mean accuracy of cross-quadrant stimulus category classifiers trained on DMC (dark grey) or PV (light grey) data for all neurons in MT (left), LIP (middle), and SC (right). Shading indicates s.d. across bootstraps. Because classifiers were trained on directions from two quadrants (one from each category) and tested on the other two quadrants (see *Methods*), classifier performance can be below chance if there is strong direction encoding in the population. **b,** Accuracy of category classifiers trained on whole-epoch data (from 200 to 400 ms after stimulus onset). **c,** same as *b* but for subpopulations of neurons without significant direction tuning indices in each area (MT: $N = 245$, LIP: $N = 352$, SC: $N = 395$). *** indicates $P < .001$. Horizontal dashed lines indicate chance level.

gorization task in which they had to correctly group motion stimuli into categories based on an arbitrarily rule in order to receive a reward, and (2) a passive viewing task in which

they viewed the same motion stimuli as for the categorization task but only had to maintain fixation on a central cue in order to receive a reward. Monkeys performed both tasks in the same session, and we could therefore compare how behavioral context affects responses to the visual stimuli during the two tasks.

We compared firing rates of individual neurons between the DMC and PV tasks, and observed that task context significantly modulates response rates in all three brain areas. We observed significantly lower firing rates for the PV task compared to the DMC task for neurons' preferred motion direction, but significantly higher firing rates for the PV task vs. the DMC task for neurons' least preferred directions. Additionally, we observed significantly weaker direction encoding during the PV task than the DMC task, both on the single-neuron and neural population level. The observation of enhanced firing rates for the preferred direction in the DMC task and increased direction encoding in the DMC task is consistent with a large body of work showing that spatial attention enhances neuronal responses and throughout the brain (Reynolds and Chelazzi, 2004; Yantis and Serences, 2003; Maunsell and Cook, 2002; Martinez-Trujillo and Treue, 2004; Treue and Maunsell, 1996; Moran and Desimone, 1985; Ignashchenkova et al., 2004; Zénon and Krauzlis, 2012). This attention-related modulation of response gain (as well as other attention-related modifications in responses, such as reductions in noise correlations between neighboring neurons) has been hypothesized to increase the discriminability of behaviorally-relevant information (Cohen and Maunsell, 2009; Cohen and Newsome, 2008). In our study, monkeys are motivated attend to the stimuli during the DMC in order to receive a reward, but do not need to actively attend to stimuli during the PV task. The difference between our experiment and previous work on spatial attention is that in the latter, there are typically multiple competing stimuli, while our tasks presented a single stimulus at a time.

Neurons in MT were not significantly category-tuned even during either the DMC task or the PV task, and DMC-PV tuning curve correlations were significantly higher in MT than in

LIP or SC. There results are consistent with previous studies showing that motion direction selectivity of MT neurons does not change as a result of perceptual or category learning (e.g., (Law and Gold, 2008; Freedman and Assad, 2006)) and suggests that neuronal characteristics of MT are not altered (or are minimally altered) by extensive training. However, MT neurons had significantly higher quadrant tuning in the DMC vs. the PV task, indicating that selectivity of MT neurons may shift slightly to accommodate current behavioral demands, although additional controls are needed to verify these results. This task-related change in single-neuron responses is compatible with work showing that stimulus context modulates patterns of noise correlations between pairs of neurons in MT (Cohen and Newsome, 2008).

Importantly, we observed markedly reduced category tuning during the PV task relative to the DMC task in LIP and SC, especially on the neural population level. It was unknown before this experiment whether category tuning in either brain area would persist even when the animals were not required to categorize the stimuli. In LIP, task experience and training can produce considerable changes in neural activity. For example, one study recorded from LIP neurons before and after monkeys learned to perform a motion categorization task (Sarma et al., 2016). Prior to training on the categorization task, monkeys performed a delayed match-to-sample task in which they compared the direction of two sequentially presented motion stimuli separated by a delay period. During this task, LIP neurons encoded the direction of the stimuli during the visual epoch (i.e., while the stimulus was bring presented), but not during the memory delay period. Monkeys were then trained on a delayed match-to-category task that used the same motion stimuli. After training, a high percentage of LIP neurons responded in a binary-like manner to motion categories, and, importantly, now showed sustained task-related delay period activity, indicating that category learning can profound change the representations within the LIP. The animals alternated between blocks of the PV and DMC tasks multiple times on each session. Every time they switched between tasks, the responses of neurons in LIP and SC flexibly shifted. These results indicate

that the neural networks that the LIP and SC are embedded in can flexibly route sensory input based on current behavioral demands.

This study has several limitations. The first limitation is that the trial structure of the PV and DMC tasks is very different: in the PV task, monkeys viewed each stimulus for 400 ms, and viewed up to 5 stimuli per trial, while in the DMC task, the sample stimulus was presented for 550 ms and was followed by a delay period during which the monkeys had remember information about the sample stimulus. Previous work has shown that neural populations in LIP form stable fixed-point attractors during a categorization task with a delay period, but not in a categorization task without a delay period (Mohan et al., 2021). The attractors compressed neural responses into a more binary-like format during the visual epoch, even before the start of the delay period, indicating that the presence of a working memory period can affect neural responses even while the stimulus is still present. Therefore, a more ideal comparison of DMC vs. PV data would have identical trial structures for the two tasks.

Another limitation of this task is that the PV task did not require animals to attend to the motion stimulus, making it difficult to distinguish between task context-related differences in neural responses from those related to the differential attentional engagement between the two tasks. One to improve on this limitation would be to require animals to perform a task on the motion stimulus that is orthogonal to the category computation. For example, animals could be trained to detect changes in luminance or motion speed of the stimuli during the PV task, although this can also introduce additional confounds in which animals attend to different features of the stimulus in the two task.

3.4 Methods

3.4.1 Subjects

Two adult (13-15 years old) male rhesus macaques (Macaca mulatta) participated in the experiment (Monkey N: ∼12 kg, Monkey S: ∼13 kg). All procedures were in accordance with the University of Chicago Institutional Animal Care and Use Committee and the National Institutes of Health guidelines and policies.

3.4.2 Behavioral tasks

For the behavioral tasks described below, the monkeys were head restrained and seated in a primate chair inserted inside an isolation box (Crist Instrument), facing a 24-inch LCD monitor on which stimuli were presented (1,920 x 1,080 resolution, refresh rate 60 Hz, 57 cm viewing distance). Reward delivery, stimulus presentation, behavioral signals, and task events were controlled by MonkeyLogic software (Asaad et al., 2013), running under MATLAB on a Windows-based PC. Gaze position was measured with an optical eye tracker (Eyelink 1000; SR Research, Ottowa, Canada) with a 1.0 kHz sample rate. For both tasks, monkeys initiated trials by holding a manual touch bar.

Delayed match-to-category task

We trained monkeys to perform a delayed match-to-category (DMC) task in which they grouped twelve dot-motion stimuli into two categories based on two orthogonal boundaries, such that motion directions that are 180° apart belong to the same category. Motion directions were separated into quadrants with three directions per quadrant, and stimuli within the same quadrant were 22.5° apart and near-boundary directions were 22.5° away from the boundary. The stimuli were 6°-diameter circular patches of white dots moving at a speed of 10°/s with 100% coherence, presented at 6.5-7.5° eccentricity in the contralateral visual

field. Animals were required to fixate within a 2.5-3.5° window.

Passive viewing paradigm

Monkeys also performed a passive viewing (PV) paradigm in which they passively viewed the same motion stimuli used in the DMC task (and presented at the same peripheral location). At the start of a trial, a blue fixation circle appeared to indicate a PV trial to monkeys (as opposed to a white fixation circle for the DMC task). Monkeys had to maintain gaze fixation on this central cue for 500 ms, after which 3-5 motion stimuli appeared in succession for 400 ms each and separated by 200 ms. Monkeys received a fluid reward at the end of the trial for maintaining gaze fixation on the central cue.

3.4.3 Surgical procedures and electrophysiological recordings

Monkeys were implanted with a titanium headpost and a single recording chamber positioned over MT, LIP and SC. Stereotaxic coordinates for chamber placement were determined from magnetic resonance imaging (MRI) scans obtained before implantation of recording chambers. MT, LIP and SC recordings were conducted in separate sessions, typically using 16- and 24-channel linear Plexon V-probes (in which channels span 1.5-2.0 mm of tissue), a dura-piercing guide tube, and a NAN microdrive system (NAN Instruments). A small subset of recording sessions from one monkey were conducted using single epoxy-insulated tungsten electrodes (FHC, Inc). We used anatomical landmarks and responses during the MGS task to guide recordings (for LIP and SC recordings). For SC recordings, we primarily targeted neurons in superficial and intermediate layers, although we also recorded neurons in deep layers as well due to the ~2mm span of recording channels on our probes. Neurophysiological signals were amplified, digitized, and stored for offline spike sorting (Plexon) to verify the quality and stability of neuronal isolation.

3.4.4 Behavioral inclusion criteria

For electrophysiological recordings and inactivation experiments, we included sessions in which behavioral performance on each category for the DMC task was at least 75% (criterion applied only to control blocks for the inactivation experiments). We excluded six LIP recording sessions (four in Monkey N and two in Monkey S) and four MT sessions (two in Monkey N and two in Monkey S) from analyses due to poor behavioral performance.

For neural analyses, we included well-isolated neurons for which we had data recorded during at least five correct trials for each of the twelve sample stimulus directions during both the DMC and PV tasks. We analyzed spiking data from 434 MT neurons recorded over 19 recording sessions (Monkey N: N neurons = 215, N sessions = 15; Monkey S: N neurons = 219, N sessions = 14), 496 LIP neurons recorded over 43 recording sessions (Monkey N: N neurons = 169, N sessions = 30; Monkey S: N neurons = 327, N sessions = 13) and 539 SC neurons recorded over 33 recording sessions (Monkey N: N neurons = 312, N sessions = 21; Monkey S: N neurons = 227, N sessions = 11).

3.4.5 Data analysis

All analyses were performed in Python v3.7.3. Behavioral analyses for the DMC task were performed on all completed trials (i.e., correct trials, misses on match trials, and false alarms on non-match trials), and neural analyses for the DMC task were performed only on correct trials. To maximize the number of repetitions per motion direction for the PV task, we included all completed trials as well as completed stimulus presentations during aborted trials (i.e., when the monkey failed to maintain fixation until the end of the trial). All P-values are two-tailed unless otherwise specified. For sliding window decoder analyses, spike trains for each neuron were smoothed using Gaussian kernel (σ = 30 ms).

Behavioral performance

To compare differences in behavioral performance on the DMC task on match trials in which the sample and test stimuli were in the same vs. opposite quadrants (**Extended Fig. 3.9b**), we computed the difference in mean accuracy for same vs. opposite quadrant match trials for each session and used a permutation test (with 5000 iterations) to randomly permute the per-session accuracy values between the two conditions.

Quantifying single-neuron direction encoding

We used three complementary methods to quantify the strength of single-neuron direction tuning or encoding. All three methods were applied on spike counts during the visual epoch from 50 to 400 ms after stimulus onset.

ROC-based direction tuning

We quantified the strength of single-neuron direction encoding using a receiver operating characteristic (ROC)-based measure. For each neuron, we applied ROC analysis to distributions of trial-by-trial epoch-average firing rates for each of the twelve stimulus conditions vs. all other conditions. We repeated this procedure 100 times for each stimulus condition, each time randomly sampling (with replacement) 10 trials per comparison, and obtained a mean rectified ROC per condition by averaging the 100 bootstrapped ROC values.

For each neuron, we computed the mean and the maximum ROC across conditions for the DMC and PV tasks and tested for significant ROC direction encoding (independently for the DMC and PV tasks) by comparing the neuron's measured ROC to a distribution of 5000 null ROC values generated by permuting the direction labels assigned to each trial. We compared each neuron's mean and maximum ROC between the DMC and PV tasks by comparing the observed difference between tasks ($\mathrm{ROC_{DMC}}$-$\mathrm{ROC_{PV}}$) to a null distribution of ROC difference values. To obtain the null distribution, we permuted the task labels

assigned to each trial and re-computed the difference in mean and max ROC between the two resulting shuffled groups. We repeated this procedure 1000 times per neuron.

Mutual information

Mutual information (MI) was calculated for each neuron from the sum of spike counts for each trial across the twelve stimulus directions during each task. The MI (in bits per spike) between spike count k and stimulus direction θ is calculated as

$$I(k;\theta) = \sum_{\theta} p(\theta) \sum_{k} p(k|\theta) \log_2 \frac{p(k|\theta)}{p(k)} \tag{3.1}$$

We tested each neuron for significant MI (independently for the DMC and PV tasks) by comparing the neuron's measured MI to a distribution of 5000 null MI values generated by permuting the direction labels assigned to each trial. We tested each neuron for differences in MI between the DMC and PV tasks by comparing the observed difference in MI between tasks ($\mathrm{MI_{DMC}}$-$\mathrm{ROC_{PV}}$) to a null distribution of MI difference values obtained by permuting the task labels assigned to each trial. We repeated this procedure 1000 times per neuron.

Direction tuning index

We computed a direction tuning index (DTI) for each neuron using the circular variance method introduced in (Mazurek et al., 2014). We calculated neurons' mean firing rate for each sample stimulus direction in a direction vector space, and quantified DTI as the normalized length of the sum of these vectors:

$$DTI = \left| \frac{\sum_{k=1}^{12} f(\theta_k) e^{i\theta_k}}{\sum_{k=1}^{12} f(\theta_k)} \right| \tag{3.2}$$

where $\mathrm{f}(\theta_\mathrm{k})$ is a neuron's mean firing rate for direction θ_k.

DTI ranges from 0 (no direction tuning) to 1. Neurons that have unimodal, cosine-like response tuning curves have DTI values close to 1, while neurons that have either uniform tuning curves or bimodal tuning curves with similar responses to directions that are 180° apart (i.e., category-tuned neurons during the DMC task) have DTI values close to 0. See **Extended Fig. 3.13** for example simulated tuning curves and the resulting DTI.

We tested each neuron for significant DTI tuning by comparing the measured DTI to a distribution of 5000 null DTIs generated by randomly shuffling the direction labels assigned to each trial. We visually verified the tuning curves of significantly direction-tuned and untuned neurons identified using this method. To further verify the validity of DTI for identifying direction-tuned neurons, we trained support vector machine (SVM) classifiers (see below) to decode stimulus direction separately on populations of neurons with significant direction tuning (during either the DMC or PV task) and on populations of neurons without significant direction tuning. In all three brain areas, performance was significantly higher for classifiers trained on DMC task activity from direction-tuned than untuned neurons (**Extended Fig. 3.14**, filled outlines). We also observed significantly higher performance on classifiers trained on PV activity from direction-tuned than untuned neurons in MT and LIP (**Extended Fig. 3.14**, unfilled outlines). To test for significant differences in DTI between the DMC and PV tasks for each neuron, we compared the observed difference (DTI_{DMC}-ROC_{DTI}) to a null distribution of DTI difference values obtained by the task label assigned to each trial and re-computed the difference in DTI between the two resulting shuffled groups. We repeated this procedure 1000 times per neuron.

For each neuron, we computed the difference between the DMC and PV data for each of these measures (i.e., ROC_{DMC}-ROC_{PV}, MI_{DMC}-MI_{PV}, and DTI_{DMC}-DTI_{PV}. We tested for significant differences between the DMC and PV tasks across neurons in each brain area by comparing, for each measure, the median of the distribution of task difference values to a distribution of 5000 null difference values, computed by randomly shuffling task labels

assigned to DMC and PV values for each neuron. **Extended Table 3.1** shows the percentage of neurons in each brain area with significantly above-chance DTI, direction ROC, and MI values, as well as the percentage of neurons with significantly higher direction-encoding values in the DMC vs. PV task.

Quantifying single-neuron category tuning

We quantified the strength and reliability of single-neuron category tuning using a receiver operating characteristic (ROC)-based category tuning index (rCTI) (Rishel et al., 2013). For each neuron, we applied ROC analysis to distributions of trial-by-trial firing rates (average across the epoch from 200 to 400 ms after stimulus onset) and compared area under the ROC curve (AUC) values for eight pairs of sample motion directions that are in the same category (Within-Category; WC) and eight pairs of directions that are in different categories (Between-Category; BC). To ensure equalized angle differences between WC and BC pairs (and thus prevent direction tuning from contaminating rCTI), the WC and BC groups each included four direction pairs spaced 45° apart and four direction pairs spaced 135° apart (**Extended Fig. 2.6a**). We quantified rCTI as the mean rectified WC AUC subtracted from the mean rectified BC AUC:

$$\begin{aligned} rCTI = \tfrac{1}{8}\sum_{p=1}^{8} 0.5 + \Big|0.5 - AUC(BC_{p1}, BC_{p2})\Big| \\ - \tfrac{1}{8}\sum_{p=1}^{8} 0.5 + \Big|0.5 - AUC(WC_{p1}, WC_{p2})\Big| \end{aligned} \tag{3.3}$$

where BC_{p1} and BC_{p2} are the two directions in the p^{th} BC pair, and WC_{p1} and WC_{p2} are the two directions in the p^{th} WC pair.

We applied the rCTI analysis to smoothed spike trains (see above) across 5-ms time steps in the trial. To generate the error shading shown in **Fig. 3.7a** we calculated rCTI for each

neuron over 500 bootstraps using 15 trials per sample motion direction (with replacement). We generated null distributions of rCTI values for each neuron using a bootstrap analysis (repeated 5000 times) in which we randomly assigned (with replacement) eight direction pairs (four 45°-spaced and four 135°-spaced pairs) to each of the shuffled BC and WC groups (**Extended Fig. 2.6**).

Support vector machine (SVM) analyses

We used SVM classifiers (with a linear kernel) to quantify the strength and timing of sample stimulus category encoding in populations of MT, LIP and SC neurons. To quantify category encoding in a direction-independent manner, we constructed cross-quadrant classifiers for which training sets consisted of trials in which the sample motion directions were from two of the four quadrants (one from each category), and testing sets consisted of sample motion direction trials from the other two quadrants (**Fig. 2.2g**). The training and testing quadrants were randomly chosen on each iteration. The analysis was applied in 5-ms steps across time in the trial and repeated for 200 iterations. For each neuron, we included 15 trials from each of six sample motion directions for training (as described above) and 15 trials from each of the remaining six sample motion directions for testing. To reduce the biases in classifier performance across brain areas due to an unequal number of neurons, for each iteration of the analysis, we randomly selected N neurons for inclusion, where N is the number of neurons in the brain area with the lower number of neurons. We generated null distributions of decoder performance values at each time using a permutation procedure (repeated 1000 times) in which we shuffled the sample direction label assigned to each trial.

We also used linear SVM classifiers to decode sample direction from MT, LIP and SC population activity. The classifiers were trained on 96 trials (8 trials from each of the 12 directions stimulus directions) and validated on 24 held-out trials (2 trials from each of the six motion directions). This analysis was applied in 5-ms steps across the trial and repeated

for 200 iterations.

3.5 Supplemental Figures and Tables

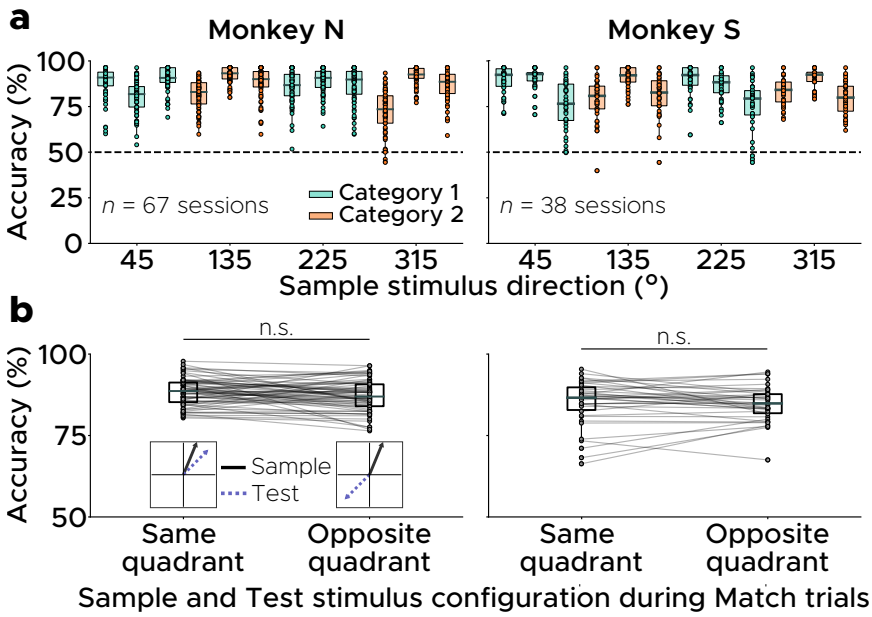

Figure 3.9: Behavioral accuracy during DMC task. a, Behavioral performance across recording sessions for each of the 12 sample stimulus directions for Monkey N (left) and Monkey S (right). Horizontal dashed lines indicate chance performance. **b,** Behavioral performance across sessions on Match trials in which the sample and test stimuli were in the same or opposite quadrants. There was no significant difference in mean performance on trials in which sample and test stimuli were in the same or different quadrants (Monkey N: Same quad. = 88.30 ± 4.21%, Opp. quad. = 87.06 ± 4.85%, $P = .124$, Monkey S: Same quad. = 85.08 ± 6.85%, Opp. quad. = 84.66 ± 5.29%, $P = .778$, permutation test.

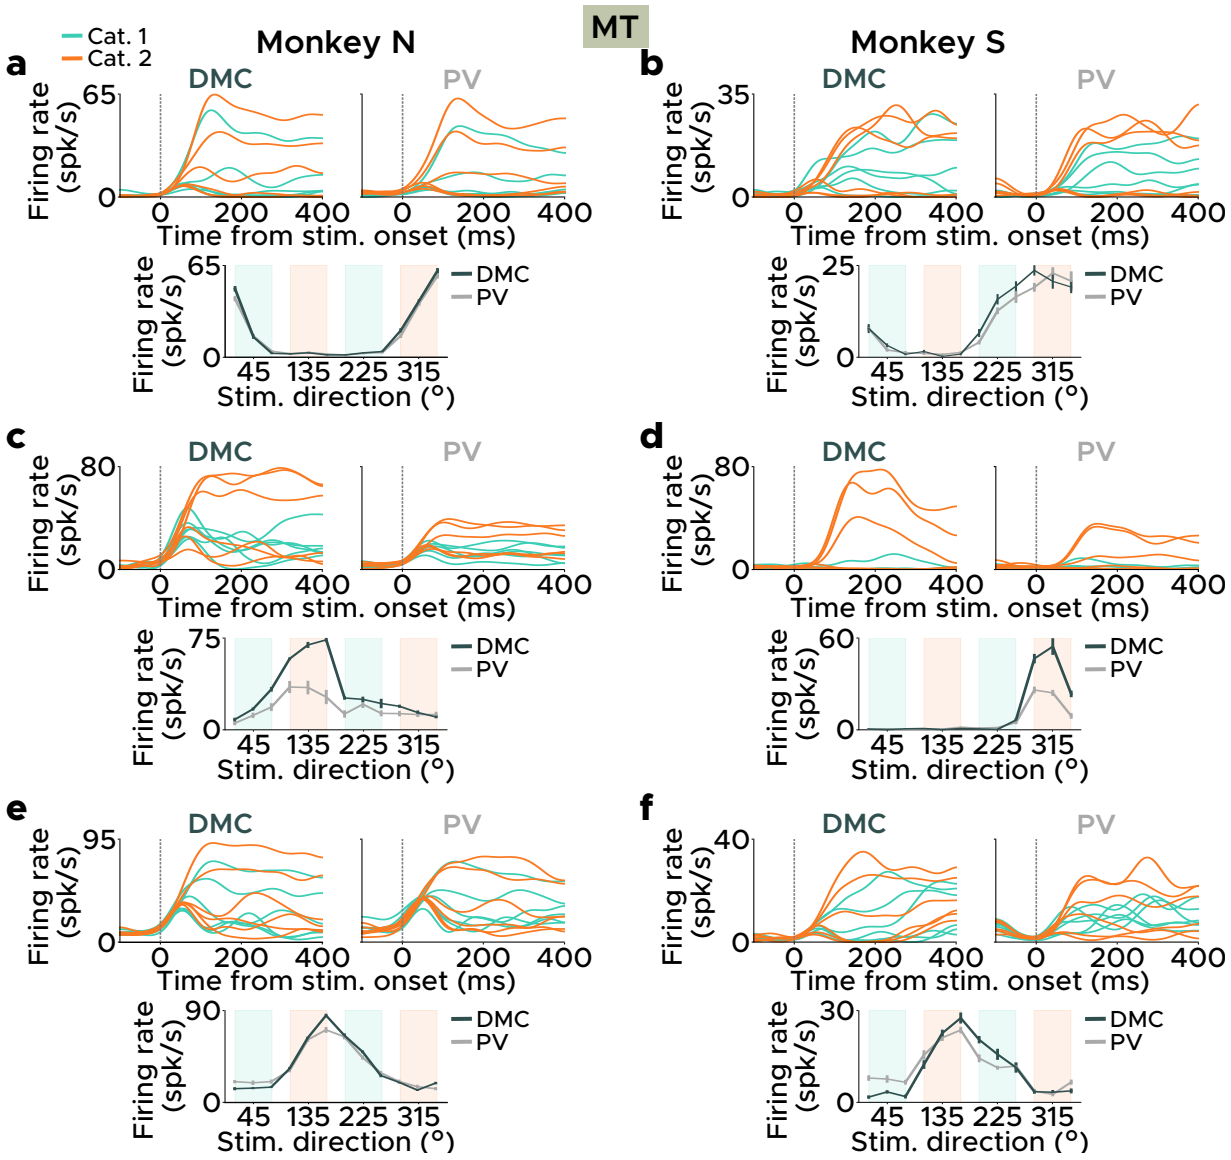

Figure 3.10: Examples of direction-tuned MT neurons that are modulated by task context. **a,** Example MT neuron recorded in Monkey N. *Left*: Peri-stimulus time histograms during the DMC (top) and PV (bottom) tasks. Color indicates stimulus category. Dotted grey lines indicate stimulus onset time. *Right*: Tuning curves showing mean firing rate (from 200-400 ms after stimulus onset) for each stimulus direction during the DMC (dark grey) and PV (light grey) tasks. Background color indicates stimulus category. Error bars indicate s.e.m. across trials. **b,** Same as *a* but for a neuron recorded in Monkey S. **c, e** Same as *a* but for two additional example MT neurons recorded in Monkey N. **d, f** Same as *b* but for two additional example MT neurons recorded in Monkey S.

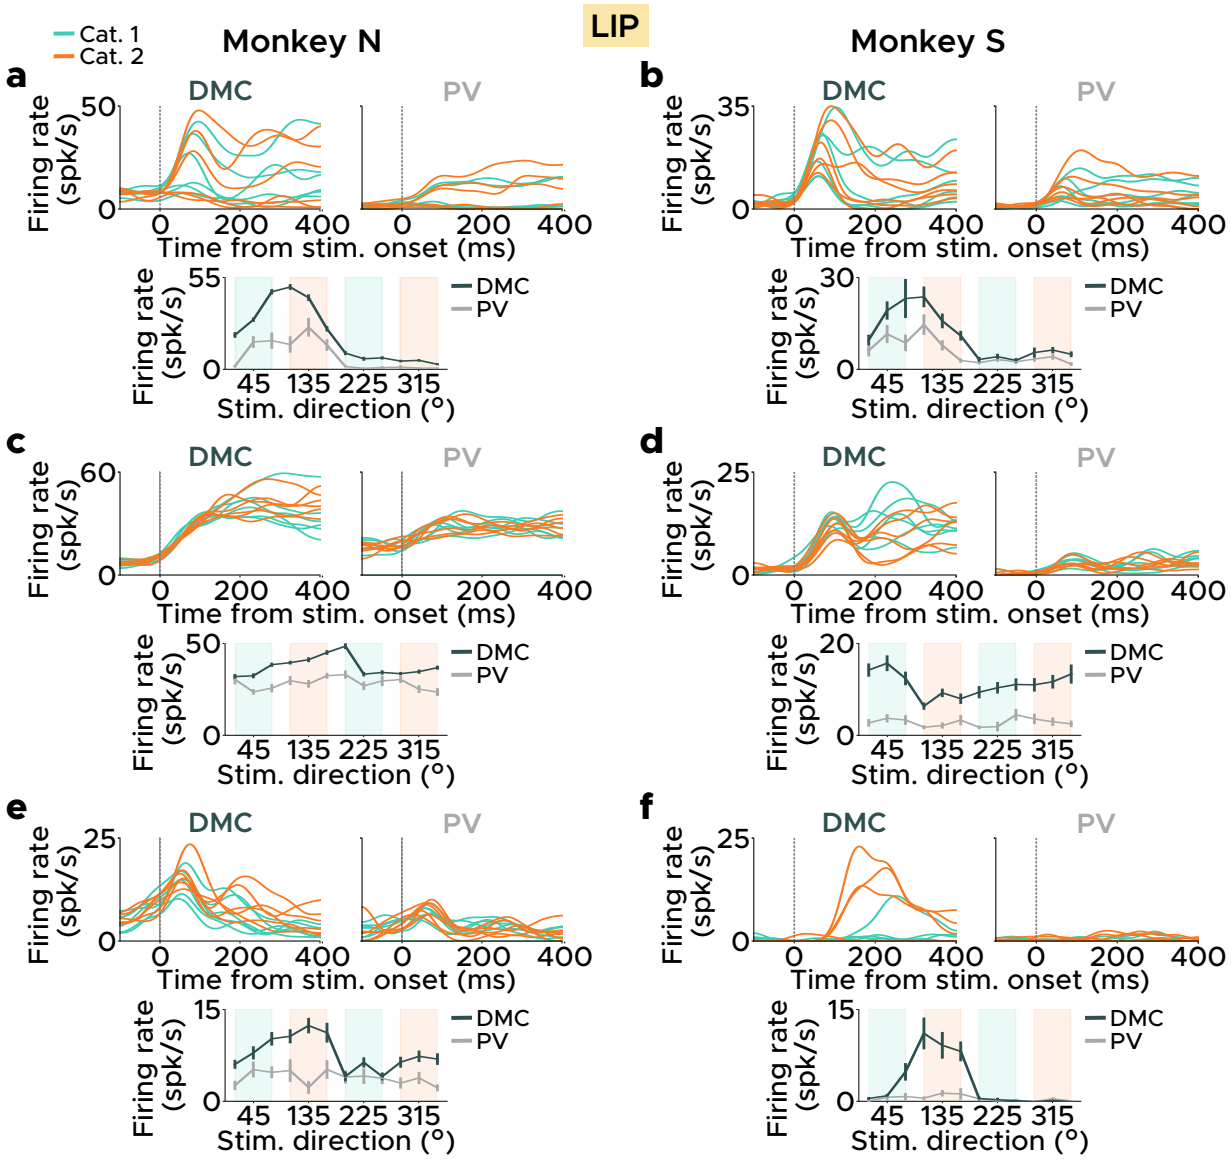

Figure 3.11: Examples of direction-tuned LIP neurons that are modulated by task context. **a,** Example LIP neuron recorded in Monkey N. *Left*: Peri-stimulus time histograms during the DMC (top) and PV (bottom) tasks. Color indicates stimulus category. Dotted grey lines indicate stimulus onset time. *Right*: Tuning curves showing mean firing rate (from 200-400 ms after stimulus onset) for each stimulus direction during the DMC (dark grey) and PV (light grey) tasks. Background color indicates stimulus category. Error bars indicate s.e.m. across trials. **b,** Same as *a* but for a neuron recorded in Monkey S. **c, e** Same as *a* but for two additional example LIP neurons recorded in Monkey N. **d, f** Same as *b* but for two additional example LIP neurons recorded in Monkey S.

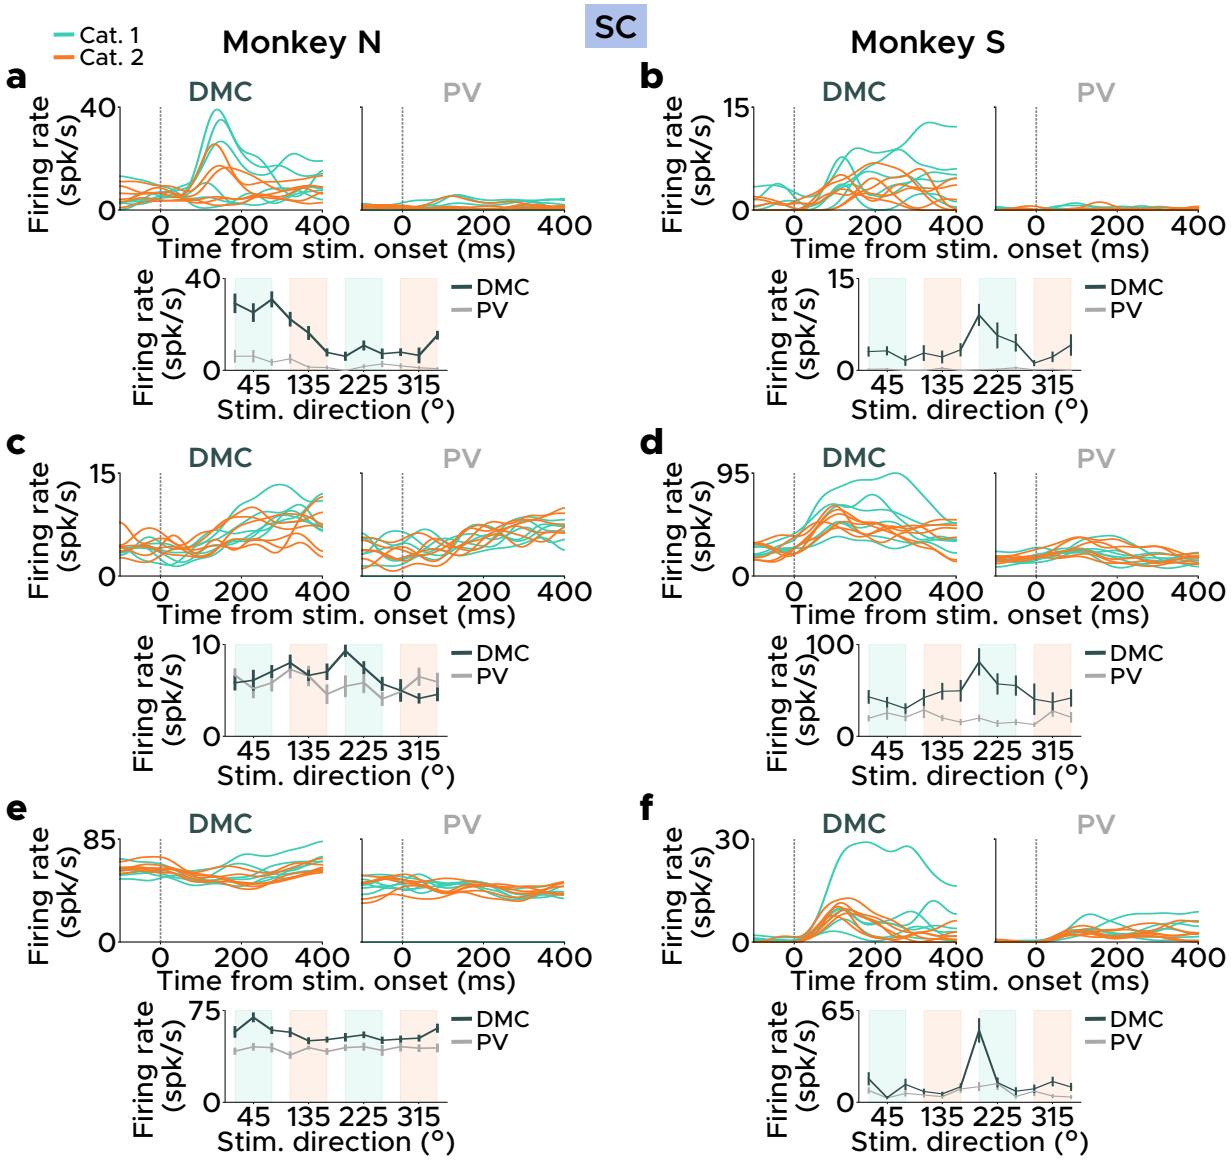

Figure 3.12: Examples of direction-tuned SC neurons that are modulated by task context. a, Example SC neuron recorded in Monkey N. *Left*: Peri-stimulus time histograms during the DMC (top) and PV (bottom) tasks. Color indicates stimulus category. Dotted grey lines indicate stimulus onset time. *Right*: Tuning curves showing mean firing rate (from 200-400 ms after stimulus onset) for each stimulus direction during the DMC (dark grey) and PV (light grey) tasks. Background color indicates stimulus category. Error bars indicate s.e.m. across trials. **b,** Same as *a* but for a neuron recorded in Monkey S. **c, e** Same as *a* but for two additional example SC neurons recorded in Monkey N. **d, f** Same as *b* but for two additional example SC neurons recorded in Monkey S.

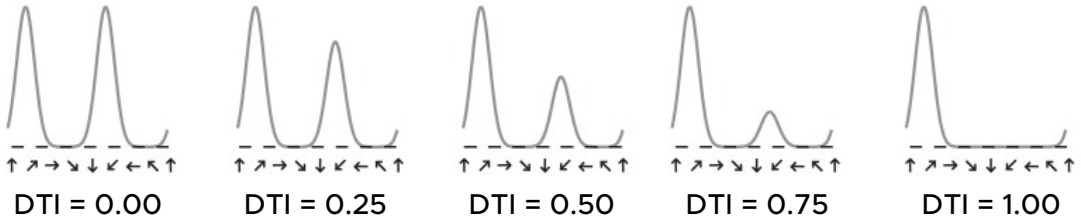

Figure 3.13: Example tuning curves and resulting direction tuning index (DTI) values. Five simulated neurons with different direction tuning properties. Each panel shows the neuron's simulated response (top) for each of the motion directions shown below. The resulting direction tuning index (DTI) value is shown below each panel. Neurons that would be category-tuned during the DMC task (i.e., respond strongly to motion directions in opposite quadrants) have DTI values close to zero (left-most neuron), while neurons that have classic Gaussian-like tuning response curve have DTI values close to 1. Figure from (Mazurek et al., 2014)

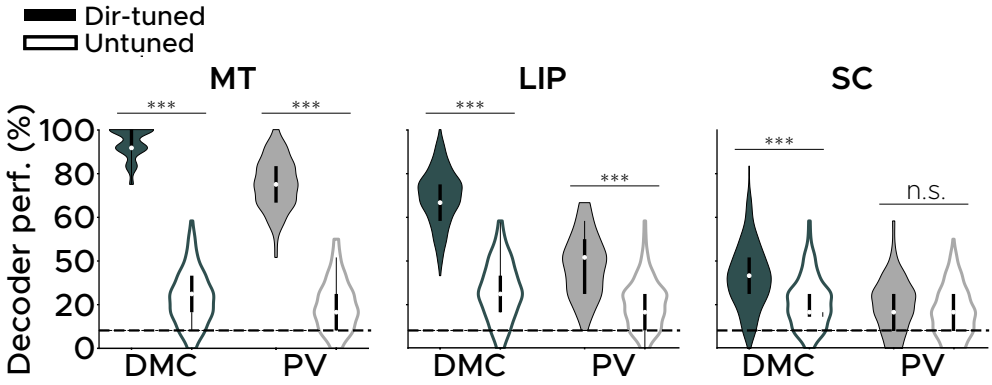

Figure 3.14: Performance of direction classifiers trained separately on direction-tuned and untuned neurons. To verify the validity of the direction tuning index (DTI) to identify direction-tuned neurons, we compared the performance of direction classifiers trained separately on neurons that had significant DTI direction tuning (during either the DMC or PV task; filled outlines, MT: $N = 189$, LIP: $N = 174$, SC: $N = 144$) and neurons without significant direction tuning during both tasks (unfilled outlines; MT: $N = 245$, LIP: $N = 352$, SC: $N = 395$). For the DMC task, performance was significantly higher for classifiers trained on tuned than untuned neurons in all three brain areas (μ ± s.d. in all areas: MT tuned = 94.25 ± 6.99%, MT untuned = 22.88 ± 13.39%; LIP tuned = 71.08 ± 12.36%, LIP untuned = 24.92 ± 12.58%; SC tuned = 30.25 ± 14.17%, SC untuned = 20.00 ± 10.37; all $P < .001$, permutation test). For the PV task, performance was significantly higher for classifiers trained on tuned vs. untuned neurons in MT and LIP (MT tuned = 73.92 ± 12.31%, MT untuned = 18.54 ± 11.97%; LIP tuned = 35.71 ± 14.0%, LIP untuned = 17.83 ± 11.52%; both $P < .001$, permutation test), but not in SC (tuned = 20.88 ± 11.58%, untuned = 15.21 ± 11.01%; $P = .266$, permutation test). Moreover, in all three brain areas, performance was significantly lower for classifiers trained on PV vs. DMC data even when only direction-tuned neurons were included (MT: $P = .002$, LIP: $P < .001$, SC: $P = .014$, permutation test). Horizontal dashed lines indicate chance performance. *** indicates $P < .001$.

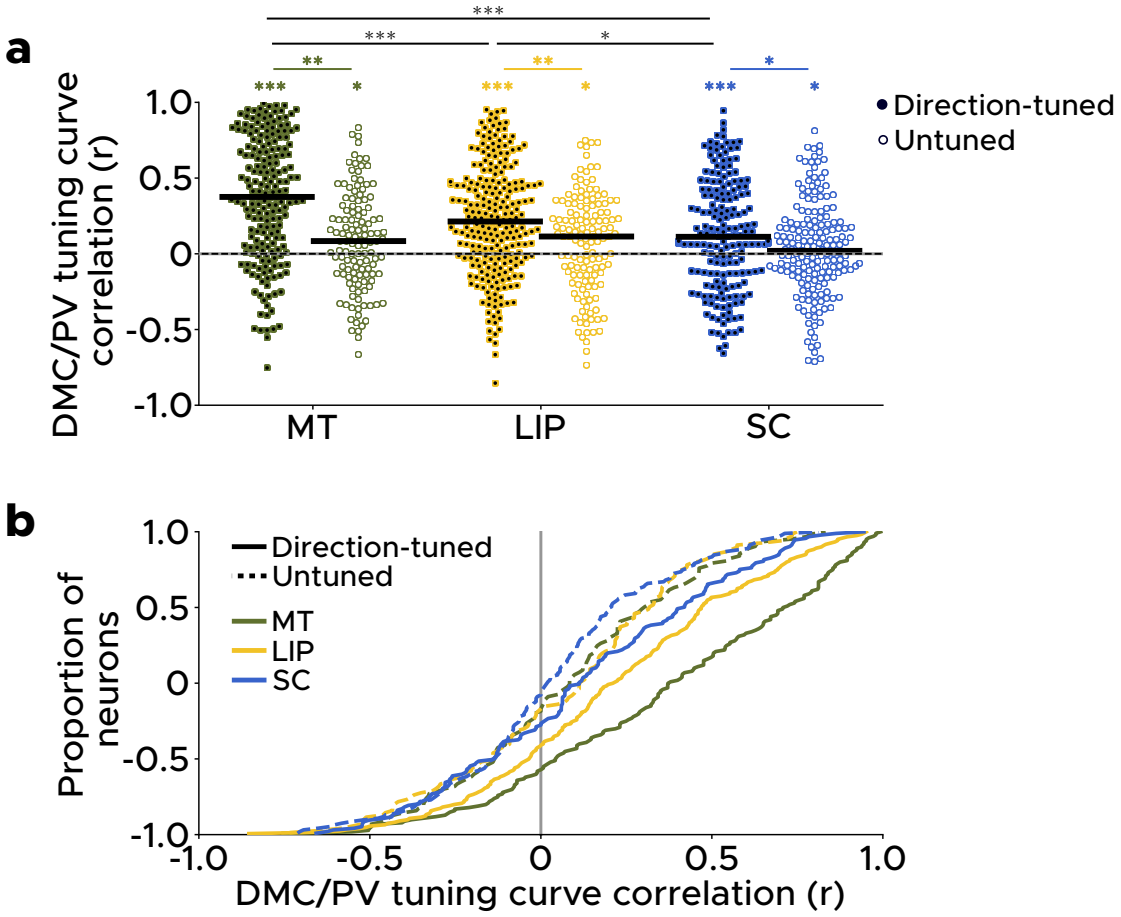

Figure 3.15: Direction-tuned neurons have more similar tuning curves during the DMC and PV tasks than untuned neurons. **a,** Distributions of tuning curve correlation r values for neurons that are significantly direction-tuned (filled circles) and untuned to direction (unfilled circles) in each area. Horizontal dashed line indicates an r value of 0. Horizontal solid black lines indicate median for each distribution. In all three brain areas, median DMC-PV tuning curve correlations were higher in direction-tuned vs. untuned neurons (MT tuned = .375 ± .493, MT untuned = .084 ± .325, P = .008, permutation test; LIP tuned = .213 ± .378 LIP untuned = .114 ± .345, P = .009; SC tuned = .112 ± .435 SC untuned = .019 ± .262, P = .018). **b,** Empirical cumulative distribution plots for direction-tuned neurons (solid traces) and untuned neurons (dashed traces) in each area. Vertical gray line indicates an r value of 0. ** indicates $P < .01$, * indicates $P < .05$.

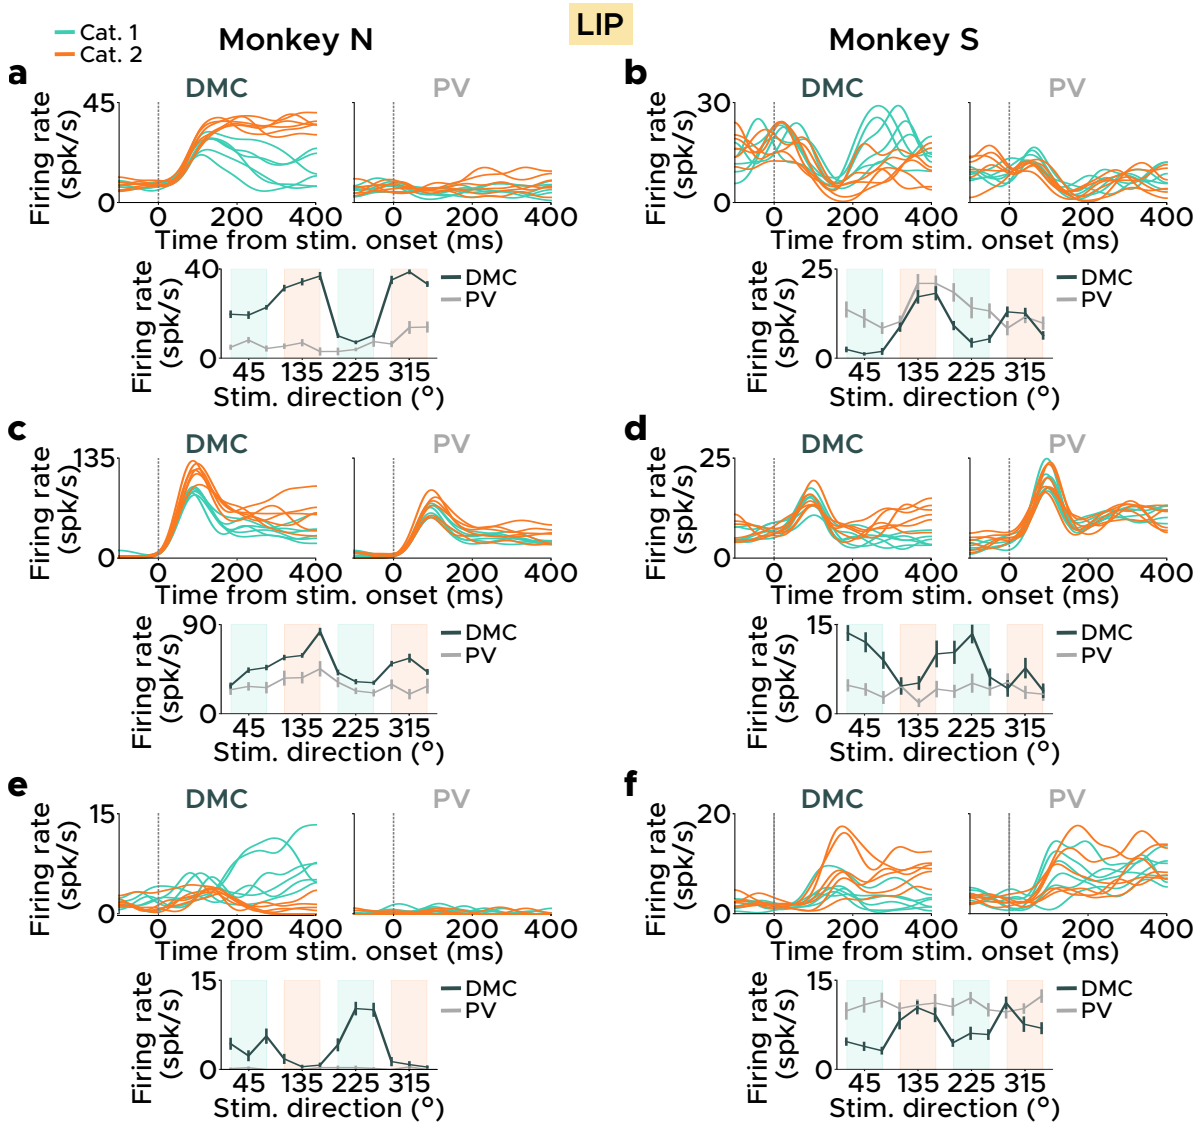

Figure 3.16: Examples of category-tuned LIP neurons that are modulated by task context. **a,** Example LIP neuron recorded in Monkey N. *Left*: Peri-stimulus time histograms during the DMC (top) and PV (bottom) tasks. Color indicates stimulus category. Dotted grey lines indicate stimulus onset time. *Right*: Tuning curves showing mean firing rate (from 200-400 ms after stimulus onset) for each stimulus direction during the DMC (dark grey) and PV (light grey) tasks. Background color indicates stimulus category. Error bars indicate s.e.m. across trials. **b,** Same as *a* but for a neuron recorded in Monkey S. **c, e** Same as *a* but for two additional example LIP neurons recorded in Monkey N. **d, f** Same as *b* but for two additional example LIP neurons recorded in Monkey S.

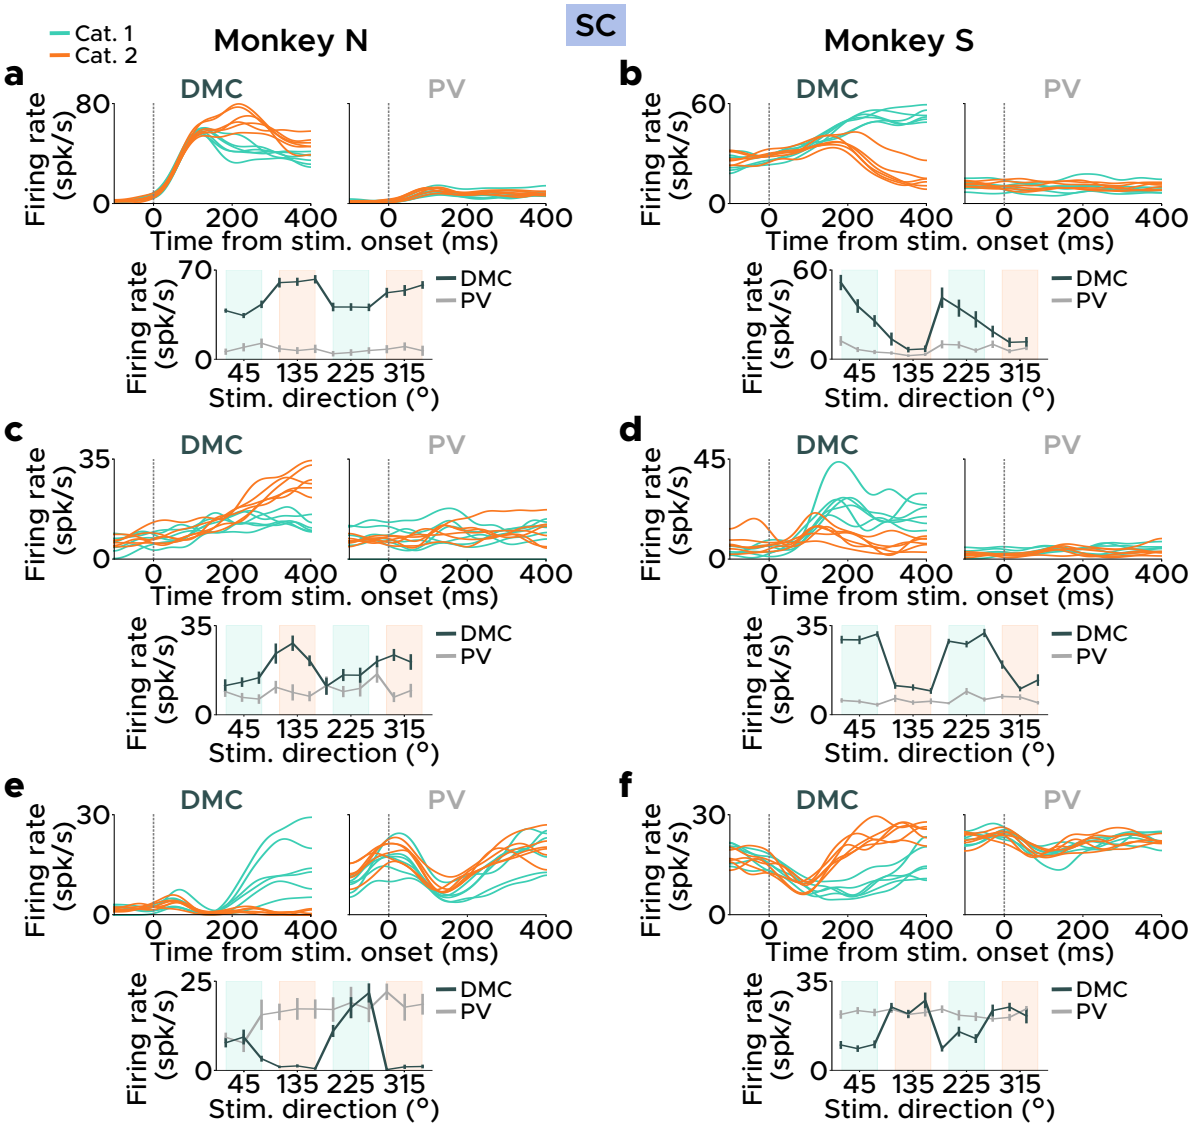

Figure 3.17: Examples of category-tuned SC neurons that are modulated by task context. a, Example SC neuron recorded in Monkey N. *Left*: Peri-stimulus time histograms during the DMC (top) and PV (bottom) tasks. Color indicates stimulus category. Dotted grey lines indicate stimulus onset time. *Right*: Tuning curves showing mean firing rate (from 200-400 ms after stimulus onset) for each stimulus direction during the DMC (dark grey) and PV (light grey) tasks. Background color indicates stimulus category. Error bars indicate s.e.m. across trials. **b,** Same as *a* but for a neuron recorded in Monkey S. **c, e** Same as *a* but for two additional example SC neurons recorded in Monkey N. **d, f** Same as *b* but for two additional example SC neurons recorded in Monkey S.

Table 3.1: Percentage of direction-encoding neurons during the DMC and PV tasks

Direction ROC								
		% DT overall		% DT by task			% Higher DTI	
Area	Monkey	DMC	PV	DMC only	PV only	DMC & PV	DMC	PV
MT	N	**54.88**	37.67	**42.55**	16.31	41.13	**60.28**	15.6
	S	**62.56**	47.95	37.5	18.45	**44.05**	**57.74**	25.0
LIP	N	**55.03**	43.79	35.65	19.13	**45.22**	**67.83**	15.65
	S	**59.63**	43.73	**41.39**	20.08	38.52	**54.1**	25.0
SC	N	**53.53**	31.09	**52.68**	18.54	28.78	**57.07**	25.37
	S	**54.63**	31.28	**53.29**	18.42	28.29	**59.87**	25.0
Mutual Information								
		% DT overall		% DT by task			% Higher DTI	
Area	Monkey	DMC	PV	DMC only	PV only	DMC & PV	DMC	PV
MT	N	**24.65**	22.33	29.41	22.06	**48.53**	**33.82**	1.47
	S	**29.22**	17.35	**49.33**	14.67	36.0	**40.0**	2.67
LIP	N	**35.5**	13.61	**67.61**	15.49	16.9	**14.08**	8.45
	S	**17.43**	13.46	**47.62**	32.14	20.24	**22.62**	3.57
SC	N	**14.1**	8.01	**59.02**	27.87	13.11	**24.59**	3.28
	S	**14.54**	11.45	**50.94**	37.74	11.32	**33.96**	5.66
Direction tuning index (DTI)								
		% DT overall		% DT by task			% Higher DTI	
Area	Monkey	DMC	PV	DMC only	PV only	DMC & PV	DMC	PV
MT	N	**35.81**	32.56	25.53	18.09	**56.38**	**15.96**	9.57
	S	**33.79**	24.66	**43.16**	22.11	34.74	**30.53**	21.05
LIP	N	**38.46**	20.71	**53.33**	13.33	33.33	**24.0**	13.33
	S	**23.24**	18.04	**40.4**	23.23	36.36	**16.16**	15.15
SC	N	13.78	**14.1**	42.11	**43.42**	14.47	**25.0**	19.74
	S	**20.26**	15.86	**47.06**	32.35	20.59	**17.65**	16.18

Percentage of neurons with significant direction encoding for each of the three direction-encoding measures (DTI, ROC, and MI). Column descriptions: "% DT overall:" % of neurons that have significant direction-tuned during the DMC and PV tasks; "% DT by task:" % of tuned neurons that are tuned only in the DMC task (DMC only), only in the PV task (PV only), or in both tasks (DMC & PV); "% Higher DTI:" % of tuned neurons that have significantly higher tuning values during the DMC or PV tasks. For each comparison in each row, the task with the higher proportion is highlighted in bold.

CHAPTER 4

DISCUSSION

4.1 Summary of results

In this thesis, we investigated the subcortical and cortical contributions to visual category processing. We trained monkeys on an abstract visual motion categorization task in which they learned to group motion directions into category based on an arbitrary category rule. Based on previous investigations of the neural mechanisms underlying categorization behaviors, the cortical lateral intraparietal area (LIP) has emerged as a leading candidate for where the direction-to-category computation may take place.

In the first chapter of this thesis, we compared the activity of the LIP and the superior colliculus (SC) during the categorization task. The SC has previously been understood to mediate spatial orienting functions, such as saccadic eye movements or covert spatial attention. We observed the superior colliculus encodes learned categories more reliably and with a shorter latency than the LIP, and that reversible inactivation of the SC causes significant deficits in animals' performance on the categorization task. In addition, we show that category and eye movement-related signals are encoded in near-orthogonal subspaces in population activity in the SC, providing an explanation for how a motor structure like the SC can be recruited to participate in more flexible cognitive behaviors. Our results reveal a novel role of the primate SC in non-spatial cognitive functions.

In the second chapter of this thesis, investigated the effects of behavioral demands on encoding of sensory and cognitive variables in the LIP, SC, and the middle temporal area (MT), an extrastriate visual area that is important for motion processing. We trained monkeys to alternate between blocks of the visual motion categorization task and a passive task in which they viewed the same motion stimuli but were not required to categorize the stimuli. In all three brain areas, the firing rates of neurons were strongly modulated by task

context, and direction encoding was significantly weaker during the PV task than the DMC task. Moreover, category encoding was largely absent in LIP and SC, indicating that neural circuits can flexibly modulate the processing of sensory information based on current task demands.

REFERENCES

Adamük E (1870) Über die Innervation der Augenbewegungen. *Zentrale Medizinische Wissenschaften* pp. 65–67.

Andersen RA, Asanuma C, Essick G, Siegel RM (1990) Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *Journal of Comparative Neurology* 296:65–113.

Antzoulatos EG, Miller EK (2011) Differences between neural activity in prefrontal cortex and striatum during learning of novel abstract categories. *Neuron* 71:243–249.

Apter JT (1946) Eye movements following strychninization of the superior colliculus of cats. *Journal of Neurophysiology* 9:73–86.

Asaad WF, Santhanam N, McClellan S, Freedman DJ (2013) High-performance execution of psychophysical tasks with complex visual stimuli in MATLAB. *Journal of Neurophysiology* 109:249–260.

Asanuma C, Andersen RA, Cowan WM (1985) The thalamic relations of the caudal inferior parietal lobule and the lateral prefrontal cortex in monkeys: Divergent cortical projections from cell clusters in the medial pulvinar nucleus. *Journal of Comparative Neurology* 241:357–381.

Basso MA, Bickford ME, Cang J (2021) Unraveling circuits of visual perception and cognition through the superior colliculus. *Neuron* 109:918–937.

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

Basso MA, Wurtz RH (1998) Modulation of neuronal activity in superior colliculus by changes in target probability. *The Journal of Neuroscience* 18:7519–7534.

Blatt GJ, Andersen RA, Stoner GR (1990) Visual receptive field organization and corticocortical connections of the lateral intraparietal area (area LIP) in the macaque. *The Journal of Comparative Neurology* 299:421–445.

Born RT, Bradley DC (2005) Structure and function of visual area mt. *Annual Review of Neuroscience* 28:157–189.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *Journal of Neuroscience* 12:4745–4765.

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

Cavanaugh J, Alvarez BD, Wurtz RH (2006) Enhanced performance with brain stimulation: attentional shift or visual cue? *The Journal of Neuroscience* 26:11347–11358.

Cavanaugh J, Wurtz RH (2004) Subcortical modulation of attention counters change blindness. *The Journal of Neuroscience* 24:11236–11243.

Clower DM, West RA, Lynch JC, Strick PL (2001) The inferior parietal lobule is the target of output from the superior colliculus, hippocampus, and cerebellum. *The Journal of Neuroscience* 21:6283–6291.

Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience* 12:1594–1600.

Cohen MR, Newsome WT (2008) Context-dependent changes in functional circuitry in visual area mt. *Neuron* 60:162–173.

Cragg BG (1969) The topography of the afferent projections in the circumstriate visual cortex of the monkey studied by the Nauta method. *Vision Research* 9:733–747.

Crapse TB, Lau H, Basso MA (2018) A role for the superior colliculus in decision criteria. *Neuron* 97:181–194.e6.

Dotson NM, Hoffman SJ, Goodell B, Gray CM (2018) Feature-based visual short-term memory is widely distributed and hierarchically organized. *Neuron* 99:215–226.e4.

Duan CA, Erlich JC, Brody CD (2015) Requirement of prefrontal and midbrain regions for rapid executive control of behavior in the rat. *Neuron* 86:1491–1503.

Duan CA, Pagan M, Piet AT, Kopec CD, Akrami A, Riordan AJ, Erlich JC, Brody CD (2021) Collicular circuits for flexible sensorimotor routing. *Nature Neuroscience* 24:1110–1120.

Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP (2016) Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications* 7:13239.

Felsen G, Mainen ZF (2008) Neural substrates of sensory-guided locomotor decisions in the rat superior colliculus. *Neuron* 60:137–148.

Ferrera VP, Yanike M, Cassanello C (2009) Frontal eye field neurons signal changes in decision criteria. *Nature Neuroscience* 12:1458–1462.

Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291:312–316.

Freedman DJ, Assad JA (2006) Experience-dependent representation of visual categories in parietal cortex. *Nature* 443:85–88.

Freedman EG, Stanford TR, Sparks DL (1996) Combined eye-head gaze shifts produced by electrical stimulation of the superior colliculus in rhesus monkeys. *Journal of Neurophysiology* 76:927–952.

Fries W (1984) Cortical projections to the superior colliculus in the macaque monkey: A retrograde study using horseradish peroxidase. *Journal of Comparative Neurology* 230:55–76.

Fritz JB (2005) Differential dynamic plasticity of a1 receptive fields during multiple spectral tasks. *Journal of Neuroscience* 25:7623–7635.

Fritz J, Shamma S, Elhilali M, Klein D (2003) Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex. *Nature Neuroscience* 6:1216–1223.

Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annual Review of Neuroscience* 30:535–574.

Hikosaka O, Wurtz RH (1985) Modification of saccadic eye movements by GABA-related substances. I. Effect of muscimol and bicuculline in monkey superior colliculus. *Journal of Neurophysiology* 53:266–291.

Hikosaka O, Wurtz RH (1986) Saccadic eye movements following injection of lidocaine into the superior colliculus. *Experimental Brain Research* 61:531–539.

Horwitz GD, Batista AP, Newsome WT (2004) Representation of an abstract perceptual decision in macaque superior colliculus. *Journal of Neurophysiology* 91:2281–2296.

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

Horwitz GD, Newsome WT (2001) Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. *Journal of Neurophysiology* 86:2543–2558.

Hussar C, Pasternak T (2010) Trial-to-trial variability of the prefrontal neurons reveals the nature of their engagement in a motion discrimination task. *Proceedings of the National Academy of Sciences* 107:21842–21847.

Hussar CR, Pasternak T (2009) Flexibility of sensory representations in prefrontal cortex depends on cell type. *Neuron* 64:730–743.

Hussar CR, Pasternak T (2012) Memory-guided sensory comparisons in the prefrontal cortex: contribution of putative pyramidal cells and interneurons. *Journal of Neuroscience* 32:2747–2761.

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

Jun EJ, Bautista AR, Nunez MD, Allen DC, Tak JH, Alvarez E, Basso MA (2021) Causal role for the primate superior colliculus in the computation of evidence for perceptual decisions. *Nature Neuroscience* 24:1121–1131.

Krauzlis RJ, Bollimunta A, Arcizet F, Wang L (2014) Attention as an effect not a cause. *Trends in Cognitive Sciences* 18:457–464.

Krauzlis RJ, Lovejoy LP, Zénon A (2013) Superior colliculus and visual spatial attention. *Annual Review of Neuroscience* 36:165–182.

Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience* 11:505–513.

Lee C, Rohrer WH, Sparks DL (1988) Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature* 332:357–360.

Lovejoy LP, Krauzlis RJ (2010) Inactivation of primate superior colliculus impairs covert selection of signals for perceptual judgments. *Nature Neuroscience* 13:261–266.

Lynch JC, Graybiel AM, Lobeck LJ (1985) The differential projection of two cytoarchitectonic subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus. *Journal of Comparative Neurology* 235:241–254.

Martinez-Trujillo JC, Treue S (2002) Attentional modulation strength in cortical area mt depends on stimulus contrast. *Neuron* 35:365–370.

Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology* 14:744–751.

Maunsell JH, van Essen DC (1983a) The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 3:2563–2586.

Maunsell JH, Van Essen DC (1983b) Functional properties of neurons in middle temporal visual area of the macaque monkey: II. Binocular interactions and sensitivity to binocular disparity. *Journal of Neurophysiology* 49:1148–1167.

Maunsell JHR, Cook EP (2002) The role of attention in visual processing. *Philosophical Transactions of the Royal Society B: Biological Sciences* 357:1063–1072.

Mazurek M, Kager M, Van Hooser SD (2014) Robust quantification of orientation selectivity and direction selectivity. *Frontiers in Neural Circuits* 8:92.

McPeek RM, Keller EL (2004) Deficits in saccade target selection after inactivation of superior colliculus. *Nature Neuroscience* 7:757–763.

Mohan K, Zhu O, Freedman DJ (2021) Interaction between neuronal encoding and population dynamics during categorization task switching in parietal cortex. *Neuron* 109:700–712.e4.

Montijn JS, Seignette K, Howlett MH, Cazemier JL, Kamermans M, Levelt CN, Heimel JA (2021) A parameter-free statistical test for neuronal responsiveness. *eLife* 10:e71969.

Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. *Science* 229:782–784.

Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK (2019) Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience* 22:1677–1686.

Mysore SP, Asadollahi A, Knudsen EI (2011) Signaling of the strongest stimulus in the owl optic tectum. *Journal of Neuroscience* 31:5186–5196.

Mysore SP, Knudsen EI (2011) Flexible categorization of relative stimulus strength by the optic tectum. *Journal of Neuroscience* 31:7745–7752.

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

Paré M, Wurtz RH (1997) Monkey posterior parietal cortex neurons antidromically activated from superior colliculus. *Journal of Neurophysiology* 78:3493–3497.

Peysakhovich B, Tetrick SM, Silva AA, Li S, Zhu O, Ibos G, Johnston WJ, Freedman DJ (2023) Primate superior colliculus is engaged in abstract higher-order cognition Pages: 2023.01.17.524416 Section: New Results.

Popovkina DV, Pasupathy A (2022) Task context modulates feature-selective responses in area v4. *Journal of Neuroscience* 42:6408–6423.

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

Ratcliff R, Hasegawa YT, Hasegawa RP, Smith PL, Segraves MA (2007) Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology* 97:1756–1774.

Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. *Annual Review of Neuroscience* 27:611–647.

Rishel CA, Huang G, Freedman DJ (2013) Independent category and spatial encoding in parietal cortex. *Neuron* 77:969–979.

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

Sarma A, Masse NY, Wang XJ, Freedman DJ (2016) Task-specific versus generalized mnemonic representations in parietal and prefrontal cortices. *Nature Neuroscience* 19:143–149.

Schiller PH, Koerner F (1971) Discharge characteristics of single units in superior colliculus of the alert rhesus monkey. *Journal of Neurophysiology* 34:920–936.

Schiller PH, Sandell JH, Maunsell JH (1987) The effect of frontal eye field and superior colliculus lesions on saccadic latencies in the rhesus monkey. *Journal of Neurophysiology* 57:1033–1049.

Schiller PH, Stryker M (1972) Single-unit recording and stimulation in superior colliculus of the alert rhesus monkey. *Journal of Neurophysiology* 35:915–924.

Scott H, Wimmer K, Pasternak T, Snyder AC (2022) Altered task demands lead to a division of labor for sensory and cognitive processing in mt. *bioRxiv* .

Shaheen LA, Slee SJ, David SV (2021) Task engagement improves neural discriminability in the auditory midbrain of the marmoset monkey. *Journal of Neuroscience* 41:284–297.

Slee SJ, David SV (2015) Rapid task-related plasticity of spectrotemporal receptive fields in the auditory midbrain. *Journal of Neuroscience* 35:13090–13102.

Sparks DL (1978) Functional properties of neurons in the monkey superior colliculus: Coupling of neuronal activity and saccade onset. *Brain Research* 156:1–16.

Stanford TR, Freedman EG, Sparks DL (1996) Site and parameters of microstimulation: evidence for independent effects on the properties of saccades evoked from the primate superior colliculus. *Journal of Neurophysiology* 76:3360–3381.

Stine GM, Trautmann EM, Jeurissen D, Shadlen MN (2022) A neural mechanism for terminating decisions. *bioRxiv* .

Swaminathan SK, Freedman DJ (2012) Preferential encoding of visual categories in parietal cortex compared to prefrontal cortex. *Nature Neuroscience* 15:315–320.

Swaminathan SK, Masse NY, Freedman DJ (2013) A comparison of lateral and medial intraparietal areas during a visual categorization task. *Journal of Neuroscience* 33:13157–13170.

Treue S, Maunsell JHR (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382:539–541.

Van Opstal AJ, Van Gisbergen JAM, Smit AC (1990) Comparison of saccades evoked by visual stimulation and collicular electrical stimulation in the alert monkey. *Experimental Brain Research* 79:299–312.

Vanegas MI, Hubbard KR, Esfandyarpour R, Noudoost B (2019) Microinjectrode system for combined drug infusion and electrophysiology. *Journal of Visualized Experiments: JoVE* p. 10.3791/60365.

Wurtz RH, Goldberg ME (1972) Activity of superior colliculus in behaving monkey. 3. Cells discharging before eye movements. *Journal of Neurophysiology* 35:575–586.

Wurtz RH, Goldberg ME (1971) Superior colliculus cell responses related to eye movements in awake monkeys. *Science* 171:82–84.

Yantis S, Serences JT (2003) Cortical mechanisms of space-based and object-based attentional control. *Current Opinion in Neurobiology* 13:187–193.

Zamarashkina P, Popovkina DV, Pasupathy A (2020) Timing of response onset and offset in macaque V4: stimulus and task dependence. *Journal of Neurophysiology* 123:2311–2325.

Zhou Y, Freedman DJ (2019) Posterior parietal cortex plays a causal role in perceptual and categorical decisions. *Science* 365:180–185.

Zhou Y, Mohan K, Freedman DJ (2022) Abstract encoding of categorical decisions in medial superior temporal and lateral intraparietal cortices. *Journal of Neuroscience* 42:9069–9081.

Zénon A, Krauzlis RJ (2012) Attention deficits without cortical neuronal deficits. *Nature* 489:434–437.