Sensorimotor transformation and information coding across cortex during perceptual decisions

by

Gerald N. Pho

B. S. E., Biomedical Engineering

Case Western Reserve University, 2010

SUBMITTED TO THE DEPARTMENT OF BRAIN AND COGNITIVE SCIENCES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN NEUROSCIENCE
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2017

Signature of Author……..……..
Department of Brain and Cognitive Sciences
April 12, 2017

Certified by……..
Mriganka Sur
Paul E. Newton Professor of Neuroscience
Thesis Supervisor

Accepted by……..……..
Matthew A. Wilson
Sherman Fairchild Professor of Neuroscience and Picower Scholar
Director of Graduate Education for Brain and Cognitive Sciences

Sensorimotor transformation and information coding across cortex during perceptual decisions

by

Gerald N. Pho

Submitted to the Department of Brain and Cognitive Sciences on
April 12, in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Neuroscience

ABSTRACT

Perceptual decision-making is an important and experimentally tractable paradigm for uncovering general principles of neural information processing and cognitive function. While the process of mapping sensory stimuli onto motor actions may appear to be simple, its neural underpinnings are poorly understood. The goal of this thesis is to better understand the neural mechanisms underlying perceptual decision-making by exploring three major questions: How is decision-relevant information encoded across the cortex? What cortical areas are necessary for perceptual decision-making? And finally, what neural mechanisms underlie the mapping of sensory percepts to appropriate motor outputs?

We investigated the roles of visual (V1), posterior parietal (PPC), and frontal motor (fMC) cortices of mice during a memory-guided visual decision task. Large-scale calcium imaging revealed that neurons in each area were heterogeneous and spanned all task epochs (stimulus, delay, response). However, information encoding was distinct across regions, with V1 encoding stimulus, fMC encoding choice, and PPC multiplexing the two variables. Optogenetic inhibition during behavior showed that all regions were necessary during the stimulus epoch, but only fMC was required during the delay and response epochs. Stimulus information was therefore rapidly transformed into behavioral choice, requiring V1, PPC, and fMC during the transformation period, but only fMC for maintaining the choice in memory prior to execution.

We further investigated whether the role of PPC was specific to visual processing or to sensorimotor transformation. Using calcium imaging during both engaged behavior and passive viewing, we found that unlike V1 neurons, most PPC neurons responded exclusively during task performance, although a minority exhibited contrast-dependent visual responses. By re-training mice on a reversed task contingency, we discovered that neurons in PPC but not V1 reflected the new sensorimotor contingency. Population analyses additionally revealed that task-specific information was represented in a dynamic code in PPC but not in V1. The strong task dependence, heterogeneity, and dynamic coding of PPC activity point to a central role in sensorimotor transformation.

By measuring and manipulating activity across multiple cortical regions, we have gained insight into how the cortex processes information during sensorimotor decisions, paving the way for future mechanistic studies using the mouse system.

Thesis Supervisor: Mriganka Sur
Title: Paul E. Newton Professor of Neuroscience

Acknowledgements

It takes a community to make a Ph.D., and mine is no exception. There are many people who made this thesis possible and who made my experience in grad school an unforgettable time of learning and personal transformation. First, I am very indebted to my thesis advisor, Mriganka Sur, who has been a wonderful mentor and support for me over these years. His contagious enthusiasm, sharp intellect, and persistent encouragement were all essential to my completion of this thesis and my development as a scientist. If given the opportunity to lead a lab, I hope that I can emulate both his inspirational passion for science as well as his unconditional generosity as a mentor. I would also like to thank the other members of my thesis committee: Bob Desimone, Chris Harvey, and Matt Wilson. Each of them played important roles over the years in providing critical feedback, helping me to write fellowships, and giving career advice.

The members of the Sur laboratory, past and present, have provided a wonderful environment to learn and do science over the years. I am fortunate to have had an amazing collaborator and mentor in Michael Goard, a former postdoc in the lab, who played a critical role not only in this thesis but also in my development as a scientist. Almost every aspect of this thesis was a team effort between the two of us, and I am grateful to have had the opportunity to work closely with someone as talented and generous as Mike. Rafiq Huda was another important collaborator and friend with whom I had countless scientific discussions that were very fruitful for my growth. And there are many others in the lab who I had the privilege of working with and learning from: Jacque Ip, Rajeev Rikhye, Ming Hu, Travis Emery, Vincent Breton-Provencher Sami El-Boustani, Vardhan Dani, Liadan Gunter, Benjamin Crawford, and Jonathan Woodson. I am especially grateful for friends like Jacque and Rajeev. We tried to “get fit” together, but much more often we just “got fat” together instead – but in the end what counts is that we got friends.

Speaking of friends, there are many, many others who not only made these years far more enjoyable for me, but who also inspired and encouraged me through all the ups and downs. The MIT Graduate Christian Fellowship was truly my family during my time here. Being part of this community not only enriched my experience of graduate school, but it also helped me to learn how to give every area of my life to Christ. I must especially thank Thomas, Sam P, Ming and Yukkee, Peng and Annie, and Sam E for inspiring me, walking with me, challenging me, and praying for me over the years. I am lucky also to have been part of the Sidney-Pacific small group for the past three years. I would especially thank Frank, Megan, Hosea, Linyi, Yunming, and AJ for being such a consistent support and source of encouragement and prayer. There are many other friends who made these years so wonderful, whether from MIT, Park Street Church or from my undergrad (Boston CWRU!): Gideon, Janet, Karen, Ben, Ruthie, Jared, Binna, Jessica, Alex, James, and Ted.

There are four friends to whom I owe a special debt of gratitude. I met three of them in my first year of grad school, and I now consider each of them closer than brothers. I am grateful for such a reliable friend in David who was a consistent voice of truth and sanity over the years. Mark, John and I have been praying together weekly for nearly six full years. While I cannot hope to describe the full impact they have had on me, I must thank them for challenging me to live out my convictions, and for lifting me up when I did not have strength to do so. And the final friend I must thank is Stephanie. Though our friendship has only blossomed relatively recently, I am tremendously grateful for the one who so patiently and persistently affirmed me over the past year and a half, and who has given me a new joy that I am hopeful will persist for years to come.

I am very proud to dedicate this thesis to my family. It has been wonderful to see my relationship with my sister JoAnn grow only deeper even as we have grown further apart in distance. I am so thankful for her encouragement and support. But most of all, I would like to thank my parents. As immigrants and refugees from Vietnam after the war, they worked tirelessly and from scratch to give me the opportunities I have today. As Chinese parents who prize discipline, integrity, and hard work, they taught me to value education and encouraged me to seek the highest possible degree I could receive. Though I can never repay them for their sacrifices for me, I am proud to represent them as the first Ph.D. in our family, and to dedicate this body of work to their unceasing love.

Last and foremost, I would like to thank and praise God, who has been my strength, my hope, and my joy over the past years and beyond. He has been faithful through all the ups and downs of my grad school years, even when I have been faithless. Looking back I stand amazed at all He has done to transform me from a self-oriented and insecure person into one with a deeper capacity to love others and pursue his beauty and truth. May this thesis be just one more testimony of all that He can and will do by the power of His love.

Table of Contents

Chapter 1: Introduction 10

1.1 Background 10

1.2 Overall organization of thesis 18

1.3 References 20

Chapter 2: Neural signatures of a sensorimotor decision in mouse visual, parietal, and frontal motor cortex 26

2.1 Summary 26

2.2 Introduction 26

2.3 Experimental Procedures 29

2.4 Results 37

2.5 Discussion 46

2.6 Contributions and Acknowledgements 50

2.7 References 51

2.8 Figures 57

Chapter 3: Causal roles of visual, parietal, and frontal motor cortex in a sensorimotor decision 65

3.1 Summary 65

3.2 Introduction 65

3.3 Experimental Procedures 70

3.4 Results 77

3.5 Discussion 86

3.6 Contributions and Acknowledgements 90

3.7 References 91

3.8 Figures 97

Chapter 4: Task-dependent representations of stimulus and choice in mouse parietal cortex 106
4.1 Summary 106
4.2 Introduction 106
4.3 Experimental Procedures 108
4.4 Results 118
4.5 Discussion 130
4.6 Contributions and Acknowledgements 135
4.7 References 135
4.8 Figures 139

Chapter 5: Dynamics of population encoding and decoding in parietal cortex during a sensorimotor decision 149
5.1 Summary 149
5.2 Introduction 149
5.3 Experimental Procedures 152
5.4 Results 161
5.5 Discussion 175
5.6 Contributions and Acknowledgements 180
5.7 References 180
5.8 Figures 186

Chapter 6: Discussion 195
6.1 Summary of key findings 195
6.2 Future directions 198
6.3 References 199

Supplementary Materials 202
Behavioral training stages 202
Video analysis of movement during delay period 202
References 203
Supplementary Figures and Tables 204

Chapter 1: Introduction

1.1 Background

1.1.1 The neural mechanisms underlying sensorimotor decisions

One of the central functions of the nervous system is to use sensory input to guide motor actions. In simple systems, deterministic sensorimotor behaviors such as the pain withdrawal reflex allow for a fast behavioral response, and are mediated by simple neural circuits that transform sensory input directly into motor action (Carew et al., 1983). However, much of cognitive behavior is non-deterministic and requires a flexible mapping of sensation onto action. Even simple perceptual decisions, such as deciding whether to go or stop at an incoming traffic light, involve a learned mapping of arbitrary sensory features (e.g. color) onto specific volitional actions (e.g. pressing the gas or the brake). Furthermore, the sensory stimulus and motor output are often separated in time, requiring short-term memory, and the mapping between them may change depending on context or experience. Such flexible, memory-guided sensorimotor transformations are thought to require more complex neural circuits that extend into the cerebral cortex (Andersen and Cui, 2009; Gold and Shadlen, 2007; Romo and de Lafuente, 2013). Understanding the neurobiology of perceptual decision-making holds great promise for generating broader insights into both the information processing principles of the brain, as well as the mechanisms of cognition in general.

The study of perceptual decision-making is very experimentally tractable, given the precise control in both magnitude and time that one can have over the inputs that lead to a decision. However, despite this apparent simplicity and decades of work in this area (Shadlen and Kiani, 2013), the field is still quite far from a comprehensive understanding of how perceptual decisions are made in the brain, and many fundamental questions remain unanswered. How is information relevant to a decision encoded and maintained in the brain? What brain areas are truly essential for perceptual decision-making? What neural mechanisms underlie the process of

mapping a sensory percept to an appropriate choice and motor output? This thesis attempts to address each of these questions by combining large-scale neural recordings, temporally-specific manipulations, and computational analysis to elucidate the neural mechanisms underlying a sensorimotor decision task.

1.1.2 Distributed encoding and maintenance of information during perceptual decisions

Perhaps the most fundamental challenge in the study of perceptual decisions (and perhaps for the field of cognitive neuroscience in general) is that the brain is a densely interconnected system, with many areas reflecting task-relevant activity even in the simplest of behaviors. The typical approach throughout the history of our field is to tame this complexity by assigning and isolating specialized functions to individual areas. In this framework, sensorimotor decisions could be performed via the (roughly) serial information processing stages of perception, cognition, and action, each implemented by distinct sets of brain regions, albeit with some feedback connections between them.

However, as intuitively appealing as this framework is, recordings of neural activity in putative sensory, cognitive, and motor regions challenge the idea of a one-to-one assignment of brain area to function (Cisek and Kalaska, 2010). For example, the maintenance of information in short-term memory is thought to be a more “cognitive” function, and indeed early work using delayed response tasks indicated the importance of higher order areas like prefrontal cortex (Funahashi et al., 1989; Fuster and Alexander, 1971; Kojima and Goldman-Rakic, 1982) and posterior parietal cortex (PPC) (Chafee and Goldman-Rakic, 1998; Constantinidis and Steinmetz, 1996; Snyder et al., 1997). However, sustained activity during the delay period has also been observed in motor cortex, subcortical regions, and even sensory areas (Nakamura and Colby, 2000; Super et al., 2001). As another important example, classic work identified signals in the

association area PPC that appear to reflect the decision process itself, as they were sensitive not only to the eventual motor choice, but also accumulated evidence for or against that decision (Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). These same “cognitive” decision signals, however, have also been observed in many other areas including putative “motor” regions like the frontal eye fields (Ding and Gold, 2012; Purcell et al., 2010), the striatum (Ding and Gold, 2010), and the superior colliculus (Horwitz and Newsome, 1999). Additionally, PPC itself has been shown to encode a panoply of other signals besides accumulated evidence, ranging from the salience of sensory stimuli (Gottlieb et al., 1998) to the direction of intended movements (Snyder et al., 1997). Therefore, although there certainly are fundamental distinctions between sensory, association, and motor areas, the highly interconnected nature of these areas and the distributed representation of task information across them makes regional assignment of individual functions a difficult challenge.

Nonetheless, identifying how various aspects of a perceptual decision task are encoded across different brain areas remains an important first step. Significant progress has been made in identifying such signals in non-human primates during perceptual decision tasks (Gold and Shadlen, 2007; Hernandez et al., 2010; Siegel et al., 2015). These studies reveal that although task-relevant information arrives first in primary sensory areas, it is subsequently distributed and reflected in the activity of many downstream areas, including frontal and parietal cortex, but also motor regions. In general, activity initially encodes information about the stimulus but is later transformed into information about the animal’s impending choice, which can be decoded not only from downstream motor circuits, but often also from sensory areas as well.

Despite this macroscopic understanding of how information flows across the brain during perceptual decisions, many mechanistic questions remain which have been difficult to address in primates. Because of the availability of powerful experimental tools, including optogenetics, two-photon microscopy, and transgenic lines, rodents, and especially mice, are quickly becoming a

model of choice for the study of sensorimotor decisions (Carandini and Churchland, 2013). Still, much work needs to be done to establish sufficiently complex decision tasks for rodents, and to measure the distributed neural signals that underlie them. In **Chapter 2** of this thesis, I present a novel, memory-guided sensorimotor decision task for mice, and demonstrate that neural signals encoding this task are distributed across sensory, association, and motor cortices.

1.1.3 Causal interrogation of brain regions underlying sensorimotor decisions

While measurement of neural activity is an important first step toward defining task-related regions, the presence of neural activity does not prove that a given region plays a causal role in mediating behavior. Causal manipulation experiments are necessary to demonstrate that task-related signals in a given brain area are not merely reflecting activity from other regions, but are in fact utilized to instruct behavior. However, these experiments occur infrequently in the decision-making literature, given that performing them in highly trained non-human primates can be prohibitively risky and costly. In addition, many perturbation experiments (whether in primates or rodents) have been inconclusive or hard to interpret for a variety of reasons, including low numbers of subjects, incomplete disruption, non-specific effects, or lack of temporal resolution.

For example, microstimulation studies have revealed that activation of neurons in sensory (Bisley et al., 2001) or parietal areas (Hanks et al., 2006) is sufficient to influence decision behavior. However, these experiments cannot distinguish whether some of the observed behavioral effects are actually due to indirect activation of other connected areas, nor can they address whether the region is necessary for task performance. Lesions can give definitive evidence for the necessity of a brain region in behavior (Gisquet-Verrier and Delatour, 2006; Sakurai and Sugimoto, 1985), but they are by nature irreversible and lack the temporal resolution to isolate specific aspects of a task. Pharmacological inactivation remains a powerful and

relatively tractable technique in both rodents (Erlich et al., 2015; Harvey et al., 2012) and primates (Katz et al., 2016; Suzuki and Gottlieb, 2013) for demonstrating whether a region is necessary for perceptual decisions. Such experiments have importantly revealed that certain areas like PPC can exhibit decision-related activity which is actually dispensable for behavioral performance in some tasks (Erlich et al., 2015; Katz et al., 2016). However, behavioral deficits are harder to interpret in these experiments, as they lack the temporal resolution to test specific task epochs, and it may thus be unclear whether a given manipulation affects stimulus perception, memory maintenance, or motor output.

The advent of optogenetics has dramatically improved our ability to manipulate neural circuits with spatial and temporal precision. Optical manipulation experiments have revealed that the behavioral effects of inactivation can depend critically on timing (Guo et al., 2014; Kopec et al., 2015; Li et al., 2016; Sachidhanandam et al., 2013). Such approaches can therefore help to identify not only what brain regions are necessary for a behavior, but also during what trial epochs is activity essential for task performance. In **Chapter 3** of this thesis, I leverage this optogenetic approach to test the necessity and temporal specificity of sensory, association, and motor cortex in a visuomotor decision task.

Clear interpretation of any perturbation experiment, however, requires an understanding of potential non-specific effects. Because the brain is a densely interconnected system, sudden manipulations of one node can have unintended effects on other nodes of the circuit. These so-called off-target effects from manipulating an area can be sufficient to alter behavioral performance, even when permanent disruptions of that area have no effect on behavior (Otchy et al., 2015). It is therefore critical to control for or measure potential off-target effects in areas that are interconnected with the brain region of interest. In **Chapter 3**, I also present experiments that involve simultaneous manipulation and recording to test whether behavioral deficits are better explained by perturbation of a single node, or by disruption of an interconnected network.

1.1.4 Sensorimotor transformation and the posterior parietal cortex

Although many areas are involved in perceptual decision-making, one brain region of keen interest to the field (and to this thesis) is the posterior parietal cortex (PPC). As an association area that receives inputs from multiple sensory areas and projects directly to motor areas, PPC has been long implicated in the process of sensorimotor transformation (Gold and Shadlen, 2007). Classic neurophysiological studies in behaving primates have shown that activity in PPC seems to reflect the decision process itself, as it encodes both the animal's impending motor action as well as the sensory evidence for that decision (Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). Such findings have been recapitulated in rodent PPC using similar accumulation of evidence decision tasks (Hanks et al., 2015).

However, the breadth and variety of responses observed in PPC makes a direct correspondence of PPC activity with accumulated evidence for a decision much too simplistic. First, PPC seems to be highly task-dependent, encoding a variety of different parameters depending on the experimental conditions, including shape selectivity (Sereno and Maunsell, 1998), categories (Freedman and Assad, 2006), visual salience (Colby and Goldberg, 1999), and even motor imagery (Aflalo et al., 2015). The diversity of PPC responses has led to vigorous debates about its specific function: whether in spatial attention (Colby and Goldberg, 1999) versus motor planning (Andersen and Buneo, 2002), or in abstract categorization (Freedman and Assad, 2011) versus action-oriented decisions (Shadlen et al., 2008). Second, even using essentially the same accumulation of evidence decision task, researchers have recently found that PPC responses can be highly heterogeneous, encoding both sensory accumulation and action selection signals (Bennur and Gold, 2011), and even multiplexing decision-irrelevant sensory signals with decision-relevant ones (Meister et al., 2013). This dizzying heterogeneity has led some to argue that PPC may consist of a dynamic network that encodes multiple signals

simultaneously, which can then be de-mixed and read-out separately by downstream structures according to behavioral needs (Park et al., 2014; Raposo et al., 2014).

Dissociating the possible roles of PPC in decisions will require approaches that go beyond mere measurement of neural activity to specific perturbation of PPC circuits during behavior (Katz et al., 2016). The enhanced experimental access offered by rodent models have led many to begin investigating the role of PPC in mice and in rats (Carandini and Churchland, 2013), yet the degree to which rodent PPC is functionally homologous to that of primates remains unclear. Both anatomical projection studies (Wang and Burkhalter, 2007; Wang et al., 2012), as well as functional mapping studies (Garrett et al., 2014; Marshel et al., 2011) indicate that the coordinates for rodent PPC overlap with a group of extrastriate areas that lie rostral to V1. Consistent with a possible role in visual processing, some have found that rat PPC is specifically necessary for visual but not auditory decisions (Licata et al., 2016; Raposo et al., 2014). Additionally, although decision signals similar to that of primates has been observed in rodent PPC (Hanks et al., 2015), inactivation experiments demonstrate that such signals are not necessary for performance of an auditory accumulation task (Erlich et al., 2015). It thus remains an open question whether rodent PPC plays a visuomotor role in a manner homologous to primate PPC, or whether it is primarily important for sensory processing as an extrastriate visual area. In **Chapter 4**, I present experiments that explicitly test these two possibilities using behavioral manipulations and calcium imaging of mouse posterior parietal cortex.

1.1.5 Population coding and dynamics with heterogeneous neurons

Great progress has been made in the early years of system neuroscience by studying the properties of single neurons in relation to externally measured variables, such as the location of a somatic sensory stimulus (Mountcastle et al., 1957) or the orientation of a visually presented

bar of light (Hubel and Wiesel, 1962). While this approach has been fruitful for early sensory regions, its success has been varied for higher cognitive areas like prefrontal cortex or PPC, even during simple decision tasks. Indeed, it has been increasingly recognized that heterogeneity and mixed selectivity across neurons is a hallmark of neural activity in these areas, and in perhaps most of the brain (Jun et al., 2010; Meister et al., 2013; Raposo et al., 2014; Rigotti et al., 2013). While the task-related responses of an individual neuron can be bewilderingly complex, population-level descriptions can often be simpler and easier to interpret (Churchland et al., 2012; Cunningham and Yu, 2014; Mante et al., 2013). Understanding the responses of one neuron with mixed selectivity may only be possible in the context of the responses of other neurons (Fusi et al., 2016), leading some to argue that the neuronal ensemble, rather than the single neuron, should be viewed as the functional unit of the nervous system (Yuste, 2015).

Information coding in heterogeneous populations can be explored using the complementary approaches of encoding and decoding. The encoding approach seeks to describe how task and stimulus information is represented in neural activity. While these encoding models have been successful in dissociating the contributions of different task-related signals to the responses of a single neuron (Chen et al., 2016; Park et al., 2014; Pillow et al., 2005; Pinto and Dan, 2015), many of these models do not take into account the ongoing activity of other neurons in the network (but see (Pillow et al., 2008; Truccolo et al., 2005)). Furthermore, merely encoding information is not enough. In order for it to be used by downstream circuits, information must be explicit enough to be decoded using simple, biologically realistic mechanisms. The decoding approach can therefore be used to test *what* type of information can be read out and *when* this information is available. But it also be used to test *how* this information is encoded across time, i.e. using which neurons. While information could be encoded in a stable manner, using the same neurons across time, others have reported that prefrontal cortex (Meyers et al., 2008; Stokes et al., 2013) and PPC (Crowe et al., 2010; Harvey et al., 2012; Morcos and Harvey, 2016) can

encode the same information at different time points using activity patterns that change dynamically (Buonomano and Maass, 2009). Whether such dynamic coding is indeed a hallmark of higher cortical regions (Murray et al., 2017) and whether it changes with task demands (Crowe et al., 2010) remains unclear. In **Chapter 5** of this thesis, I employ both encoding and decoding approaches to examine the population dynamics and coding principles of mouse PPC during perceptual decision-making.

1.2 Overall organization of thesis

This thesis can be divided into two major sections, each consisting of two chapters. Both sections are intimately related as they involve investigation of the same cortical areas in the mouse brain, during variations of the same perceptual decision task.

The first major section centers on the neural mechanisms of a memory-guided sensorimotor decision task, and the bulk of the findings have appeared in publication (Goard et al., 2016). **Chapter 2** describes the behavioral task and the task-related responses of neurons in three cortical areas, measured using population calcium imaging: primary visual cortex (V1), posterior parietal cortex (PPC), and frontal motor cortex (fMC). Individual neuron responses were heterogeneous within each area, with neurons responding during each of the three task epochs (stimulus, delay, and response) distributed across the cortex. Population analyses, however, revealed that each area dynamically encoded different types of information, with V1 encoding stimulus, fMC encoding choice, and PPC encoding stimulus early and then choice later in the trial. **Chapter 3** then describes an optogenetic approach to causally test the necessity of each of these regions for behavior with temporal specificity. While all three areas were necessary during the stimulus epoch, only fMC was necessary during the delay and response epochs. This led to a model in which stimulus information is transformed rapidly into a motor plan, which is represented in the sustained activity of motor cortex neurons until execution of the response. In addition to the

published behavioral experiments (Goard et al., 2016), Chapter 3 also describes work that probed specifically at the role of PPC and its interactions with V1. Simultaneous recording and inactivation experiments demonstrate that PPC inactivation affected V1 responses via network-level effects, suggesting that PPC may provide permissive feedback input to V1 that is necessary for task performance.

The second major section of the thesis focuses on understanding the function and information coding of the posterior parietal cortex (PPC) during flexible perceptual decision-making, work that is in preparation for publication. **Chapter 4** describes variations on the same behavioral decision task which enabled distinctions between a potential role for PPC in sensory versus sensorimotor processing. Calcium imaging experiments during engaged visual behavior and passive viewing of the same stimuli revealed that unlike neurons in V1, the majority of PPC neurons were strongly gated by task performance, although a subset of PPC neurons showed contrast-modulated visual responses. Imaging experiments performed before and after task contingency reversal further demonstrated that most PPC neurons were sensitive to the learned sensorimotor contingency, rather than to the stimulus itself. **Chapter 5** probes more deeply into the information coding properties of PPC using population-level analyses. An encoding model relating task variables to PPC responses was dramatically improved when the activity of the network was incorporated. Additionally, time-dependent decoding analyses provided evidence that PPC, unlike V1, represented task-relevant information using a dynamic code.

Finally, in **Chapter 6,** the findings of the preceding chapters are integrated and discussed in the broader context of sensorimotor decision-making and its underlying neural mechanisms.

1.3 References

Aflalo, T., Kellis, S., Klaes, C., Lee, B., Shi, Y., Pejsa, K., Shanfield, K., Hayes-Jackson, S., Aisen, M., Heck, C.*, et al.* (2015). Neurophysiology. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. Science *348*, 906-910.

Andersen, R.A., and Buneo, C.A. (2002). Intentional maps in posterior parietal cortex. Annu Rev Neurosci *25*, 189-220.

Andersen, R.A., and Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits. Neuron *63*, 568-583.

Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. J Neurosci *31*, 913-921.

Bisley, J.W., Zaksas, D., and Pasternak, T. (2001). Microstimulation of cortical area MT affects performance on a visual working memory task. J Neurophysiol *85*, 187-196.

Buonomano, D.V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. Nat Rev Neurosci *10*, 113-125.

Carandini, M., and Churchland, A.K. (2013). Probing perceptual decisions in rodents. Nat Neurosci *16*, 824-831.

Carew, T.J., Hawkins, R.D., and Kandel, E.R. (1983). Differential classical conditioning of a defensive withdrawal reflex in Aplysia californica. Science *219*, 397-400.

Chafee, M.V., and Goldman-Rakic, P.S. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. J Neurophysiol *79*, 2919-2940.

Chen, J.L., Voigt, F.F., Javadzadeh, M., Krueppel, R., and Helmchen, F. (2016). Long-range population dynamics of anatomically defined neocortical networks. Elife *5*.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. Nature *487*, 51-56.

Cisek, P., and Kalaska, J.F. (2010). Neural mechanisms for interacting with a world full of action choices. Annu Rev Neurosci *33*, 269-298.

Colby, C.L., and Goldberg, M.E. (1999). Space and attention in parietal cortex. Annu Rev Neurosci *22*, 319-349.

Constantinidis, C., and Steinmetz, M.A. (1996). Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. J Neurophysiol *76*, 1352-1355.

Crowe, D.A., Averbeck, B.B., and Chafee, M.V. (2010). Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. J Neurosci *30*, 11640-11653.

Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. Nat Neurosci *17*, 1500-1509.

Ding, L., and Gold, J.I. (2010). Caudate encodes multiple computations for perceptual decisions. The Journal of neuroscience : the official journal of the Society for Neuroscience *30*, 15747-15759.

Ding, L., and Gold, J.I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. Cereb Cortex *22*, 1052-1067.

Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. Elife *4*.

Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. Nature *443*, 85-88.

Freedman, D.J., and Assad, J.A. (2011). A proposed common neural mechanism for categorization and perceptual decisions. Nat Neurosci *14*, 143-146.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J Neurophysiol *61*, 331-349.

Fusi, S., Miller, E.K., and Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. Current opinion in neurobiology *37*, 66-74.

Fuster, J.M., and Alexander, G.E. (1971). Neuron activity related to short-term memory. Science *173*, 652-654.

Garrett, M.E., Nauhaus, I., Marshel, J.H., and Callaway, E.M. (2014). Topography and areal organization of mouse visual cortex. J Neurosci *34*, 12587-12600.

Gisquet-Verrier, P., and Delatour, B. (2006). The role of the rat prelimbic/infralimbic cortex in working memory: not involved in the short-term maintenance but in monitoring and processing functions. Neuroscience *141*, 585-596.

Goard, M.J., Pho, G.N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife *5*, 558.523.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annual review of neuroscience *30*, 535-574.

Gottlieb, J.P., Kusunoki, M., and Goldberg, M.E. (1998). The representation of visual salience in monkey parietal cortex. Nature *391*, 481-484.

Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014). Flow of cortical activity underlying a tactile decision in mice. Neuron *81*, 179-194.

Hanks, T.D., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nat Neurosci *9*, 682-689.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220-223.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62-68.

Hernandez, A., Nacher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vazquez, Y., Camarillo, L., and Romo, R. (2010). Decoding a perceptual decision process across cortex. Neuron *66*, 300-314.

Horwitz, G.D., and Newsome, W.T. (1999). Separate signals for target selection and movement specification in the superior colliculus. Science *284*, 1158-1161.

Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. The Journal of physiology *160*, 106-154.

Jun, J.K., Miller, P., Hernandez, A., Zainos, A., Lemus, L., Brody, C.D., and Romo, R. (2010). Heterogenous population coding of a short-term memory and decision task. J Neurosci *30*, 916-929.

Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. Nature *535*, 285-288.

Kojima, S., and Goldman-Rakic, P.S. (1982). Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. Brain Res *248*, 43-49.

Kopec, C.D., Erlich, J.C., Brunton, B.W., Deisseroth, K., and Brody, C.D. (2015). Cortical and Subcortical Contributions to Short-Term Memory for Orienting Movements. Neuron *88*, 367-377.

Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. Nature *532*, 459-464.

Licata, A.M., Kaufman, M.T., Raposo, D., Ryan, M.B., Sheppard, J.P., and Churchland, A.K. (2016). Posterior parietal cortex guides visual decisions in rats. bioRxiv.

Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78-84.

Marshel, J.H., Garrett, M.E., Nauhaus, I., and Callaway, E.M. (2011). Functional specialization of seven mouse visual cortical areas. Neuron *72*, 1040-1054.

Meister, M.L., Hennig, J.A., and Huk, A.C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. J Neurosci *33*, 2254-2267.

Meyers, E.M., Freedman, D.J., Kreiman, G., Miller, E.K., and Poggio, T. (2008). Dynamic population coding of category information in inferior temporal and prefrontal cortex. J Neurophysiol *100*, 1407-1419.

Morcos, A.S., and Harvey, C.D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. Nat Neurosci *19*, 1672-1681.

Mountcastle, V.B., Davies, P.W., and Berman, A.L. (1957). Response properties of neurons of cat's somatic sensory cortex to peripheral stimuli. J Neurophysiol *20*, 374-407.

Murray, J.D., Bernacchia, A., Roy, N.A., Constantinidis, C., Romo, R., and Wang, X.J. (2017). Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. Proceedings of the National Academy of Sciences of the United States of America *114*, 394-399.

Nakamura, K., and Colby, C.L. (2000). Visual, saccade-related, and cognitive activation of single neurons in monkey extrastriate area V3A. J Neurophysiol *84*, 677-692.

Otchy, T.M., Wolff, S.B., Rhee, J.Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S.M., and Olveczky, B.P. (2015). Acute off-target effects of neural circuit manipulations. Nature *528*, 358-363.

Park, I.M., Meister, M.L., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. Nat Neurosci *17*, 1395-1403.

Pillow, J.W., Paninski, L., Uzzell, V.J., Simoncelli, E.P., and Chichilnisky, E.J. (2005). Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. J Neurosci *25*, 11003-11013.

Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995-999.

Pinto, L., and Dan, Y. (2015). Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed Behavior. Neuron *87*, 437-450.

Purcell, B.A., Heitz, R.P., Cohen, J.Y., Schall, J.D., Logan, G.D., and Palmeri, T.J. (2010). Neurally constrained modeling of perceptual decision making. Psychol Rev *117*, 1113-1143.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat Neurosci *17*, 1784-1792.

Rigotti, M., Barak, O., Warden, M.R., Wang, X.J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature *497*, 585-590.

Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci *22*, 9475-9489.

Romo, R., and de Lafuente, V. (2013). Conversion of sensory signals into perceptual decisions. Progress in neurobiology *103*, 41-75.

Sachidhanandam, S., Sreenivasan, V., Kyriakatos, A., Kremer, Y., and Petersen, C.C. (2013). Membrane potential correlates of sensory perception in mouse barrel cortex. Nat Neurosci *16*, 1671-1677.

Sakurai, Y., and Sugimoto, S. (1985). Effects of lesions of prefrontal cortex and dorsomedial thalamus on delayed go/no-go alternation in rats. Behavioural brain research *17*, 213-219.

Sereno, A.B., and Maunsell, J.H. (1998). Shape selectivity in primate lateral intraparietal cortex. Nature *395*, 500-503.

Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. Neuron *80*, 791-806.

Shadlen, M.N., Kiani, R., Hanks, T.D., and Churchland, A.K. (2008). Neurobiology of Decision Making An Intentional Framework. In, C. Engel, and W. Singer, eds. (MIT Press), pp. 71-102.

Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. Journal of neurophysiology *86*, 1916-1936.

Siegel, M., Buschman, T.J., and Miller, E.K. (2015). Cortical information flow during flexible sensorimotor decisions. Science *348*, 1352-1355.

Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Coding of intention in the posterior parietal cortex. Nature *386*, 167-170.

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

Super, H., Spekreijse, H., and Lamme, V.A. (2001). A neural correlate of working memory in the monkey primary visual cortex. Science *293*, 120-124.

Suzuki, M., and Gottlieb, J. (2013). Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. Nat Neurosci *16*, 98-104.

Truccolo, W., Eden, U.T., Fellows, M.R., Donoghue, J.P., and Brown, E.N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. J Neurophysiol *93*, 1074-1089.

Wang, Q., and Burkhalter, A. (2007). Area map of mouse visual cortex. The Journal of comparative neurology *502*, 339-357.

Wang, Q., Sporns, O., and Burkhalter, A. (2012). Network analysis of corticocortical connections reveals ventral and dorsal processing streams in mouse visual cortex. J Neurosci *32*, 4386-4399.

Yuste, R. (2015). From the neuron doctrine to neural networks. Nat Rev Neurosci *16*, 487-497.

Chapter 2: Neural signatures of a sensorimotor decision in mouse visual, parietal, and frontal motor cortex[1].

2.1 Summary

Mapping specific sensory features to future motor actions is a crucial capability of mammalian nervous systems, but the neural underpinnings are not well understood. We investigated the role of visual (V1), posterior parietal (PPC), and frontal motor (fMC) cortices for sensorimotor mapping in mice during performance of a memory-guided visual discrimination task using large-scale calcium imaging. Single neuron responses within each region were surprisingly heterogeneous and spanned all task epochs (stimulus, delay, response), with a fraction of neurons in each area exhibiting sustained activity that varied parametrically with delay duration. Population analyses demonstrated unique encoding of stimulus identity and behavioral choice information across regions, with V1 encoding stimulus, fMC encoding choice even early in the trial, and PPC multiplexing the two variables. These results suggest a model in which stimulus identity is rapidly transformed into a behavioral choice and maintained in higher regions until task execution.

2.2 Introduction

The ability to use sensory input to guide motor action is a principal task of the nervous system. Simple sensorimotor transformations, such as the patellar reflex, can be mediated by simple neural circuits within the peripheral nervous system. However, more sophisticated sensorimotor decisions, like using a traffic signal to guide future driving maneuvers, often requires mapping specific sensory features to motor actions at a later time, and are thought to require more complex neural circuits extending into the cerebral cortex (Andersen and Cui, 2009; Gold and Shadlen, 2007; Romo and de Lafuente, 2013).

[1] The entirety of the findings presented in this chapter appeared in Goard, Pho et al., 2016

Over the past several decades, a number of researchers have measured neural activity during memory-guided sensorimotor decisions using delayed-response and working memory tasks. However, despite the wealth of research in this area, there are a number of unresolved questions. First, it is unclear which regions are responsible for sensorimotor transformation. For example, single-unit electrophysiological recordings (Andersen and Cui, 2009; Bennur and Gold, 2011; Freedman and Assad, 2006; Gold and Shadlen, 2007; Shadlen and Newsome, 2001) and pharmacological inactivation (Li et al., 1999) (but see (Chafee and Goldman-Rakic, 2000)) studies in non-human primates have implicated posterior parietal cortex (PPC) in mapping sensory input to appropriate motor responses. However, recent rat auditory (Erlich et al., 2015) and mouse whisker (Guo et al., 2014b) studies have challenged this view, finding no role for PPC in memory-guided sensorimotor decisions. Rather, a number of studies suggest that cortical (Guo et al., 2014b; Murakami et al., 2014; Zagha et al., 2015) and subcortical (Kopec et al., 2015; Znamenskiy and Zador, 2013) motor regions might be the key site for sensorimotor transformation. In contrast, PPC does appear to be necessary for visual sensorimotor decision tasks in mice (Harvey et al., 2012) and rats (Raposo et al., 2014); though the proximity of PPC to secondary visual regions makes it difficult to clearly isolate the effect of PPC inactivation (see discussion in (Erlich et al., 2015)). There are several possibilities for the discrepancies seen in the previous studies, including differences in species, sensory modality, and behavioral task. In order to further investigate the locus of sensorimotor transformation, we hoped in this work to take a more comprehensive approach toward measuring and perturbing activity across sensory, parietal association, and motor cortical regions during a delayed-response task.

A second unresolved issue lies in determining which region(s) are responsible for maintaining task-relevant information in the delay period between the stimulus and response. For decades, researchers have used delayed-response tasks to study the maintenance of information (short-term memory), and have observed the presence of sustained neural activity in distributed

cortical and subcortical structures; including prefrontal cortex (Funahashi et al., 1989; Fuster and Alexander, 1971; Kojima and Goldman-Rakic, 1982; Miller et al., 1996; Romo et al., 1999; Sreenivasan et al., 2014), parietal cortex (Chafee and Goldman-Rakic, 1998; Constantinidis and Steinmetz, 1996; Harvey et al., 2012; Shadlen and Newsome, 2001; Snyder et al., 1997), sensory (Nakamura and Colby, 2000; Super et al., 2001) and motor cortices (de Lafuente and Romo, 2005; di Pellegrino and Wise, 1993; Erlich et al., 2015; Guo et al., 2014b; Hernandez et al., 2010; Li et al., 2015), as well as several subcortical regions (Fuster and Alexander, 1973; Kawagoe et al., 1998). Building on this work, theoretical studies have yielded biologically-plausible models describing how sustained neural activity can be generated and maintained to support short-term memory (Compte et al., 2000; Goldman et al., 2003; Wang, 2008). However, recent studies in rodents have indicated that some of the regions traditionally thought to be crucial for short-term memory maintenance, such as parietal and prefrontal cortices, appear not to play a major role in some tasks (Erlich et al., 2015; Guo et al., 2014b; Liu et al., 2014), though other studies have found an important role for sustained activity in motor regions consistent with theoretical models (Kopec et al., 2015; Li et al., 2016). Is the sustained activity observed in parietal and prefrontal cortex epiphenomenal? Or might the difference stem from another aspect of the task (e.g., modality)? We reasoned that we could help resolve this issue by measuring and perturbing activity in distributed cortical regions during a visual delayed-response task.

To leverage the genetic tractability of the mouse toward a better mechanistic understanding of memory-guided sensorimotor decisions, we modified a visual discrimination task used for head-fixed mice (Andermann et al., 2010) by using a retractable spout to separate the components of the task into discrete stimulus, delay, and response epochs. We took advantage of high sensitivity genetically-encoded calcium indicators (Chen et al., 2013) to enable high-yield 2-photon volume scanning approach (Kampa et al., 2011). This allowed us to image from large populations of neurons in sensory, association and frontal motor regions during

discrete epochs of a memory-guided sensorimotor decision task, including stimulus encoding, delay and behavioral response.

2.3 Experimental Procedures

2.3.1 Surgical procedures

All procedures were approved by the Massachusetts Institute of Technology Animal Care and Use Committee. Data were collected from male adult (60-120 day old) wild-type mice (C57BL/6; n = 12). The animals were housed on a 12 hour light/dark cycle in cages of up to 5 animals before the implants, and individually after the implants. All surgeries were conducted under isoflurane anesthesia (3.5% induction, 1.5-2.5% maintenance). Meloxicam (1 mg kg^{-1}, subcutaneous) was administered pre-operatively and every 24 hours for 3 days to reduce inflammation. Once anesthetized, the scalp overlying the dorsal skull was sanitized and removed. The periosteum was removed with a scalpel and the skull was abraded with a drill burr to improve adhesion of dental acrylic. Stereotaxic coordinates for future viral injections were marked with a non-toxic ink and covered with a layer of silicon elastomer (Kwik-Sil, World Precision Instruments) to prevent acrylic bonding. The entire skull surface was then covered with dental acrylic (C&B-Metabond, Parkell) mixed with black ink to reduce light transmission. A custom-designed stainless steel head plate (eMachineShop.com) was then affixed using dental acrylic. After head plate implantation, mice recovered for at least five days before beginning water restriction.

After behavioral training was complete, animals were taken off water restriction for five days before undergoing a second surgery to implant the imaging window(s). Procedures for anesthetic administration and post-operative care were identical to the first surgery. The dental acrylic and silicon elastomer covering the targeted region were removed using a drill burr. The skull surface was then cleaned and a craniotomy (2-4 mm, depending on targeted structure) was

made over the region of interest, leaving the dura intact. For imaging experiments, neurons were labeled with a genetically-encoded calcium indicator by microinjection (Stoelting) of 50 nl AAV2/1.Syn.GCaMP6s.WPRE.SV40 (University of Pennsylvania Vector Core; diluted to a titer of 10^{12} genomes ml^{-1}) 300 μm below the pial surface. Between two and five injections were made in each exposed region, centered at V1 (4.2 mm posterior, 2.5 mm lateral to Bregma), PPC (2 mm posterior, 1.7 mm lateral to Bregma) or fMC (1.5 mm anterior, 1 mm lateral to Bregma). Since the viral expression spreads laterally from the injection site, exact stereotaxic locations were photographed through the surgical microscope for determining imaging areas. Finally, a cranial window was implanted over the craniotomy and sealed first with silicon elastomer then with dental acrylic. The cranial windows were made of two rounded pieces of coverglass (Warner Instruments) bonded with optical glue (NOA 61, Norland). The bottom piece was circular or oval, custom cut according to cortical region(s) (V1: 2.5 mm x 2.5 mm; PPC: 1 mm x 2 mm; V1 + PPC: 4 mm x 4 mm; fMC: 2 mm x 2.5 mm, anterior-posterior x medial-lateral) and fit snugly in the craniotomy. The top piece was a larger circular coverglass (3-5 mm, depending on size of bottom piece) and was bonded to the skull using dental acrylic. Mice recovered for five days before commencing water restriction.

2.3.2 Behavior

Mice were head-fixed using optical hardware (Thorlabs) and placed in a polypropylene tube to limit movement. Spout position was controlled by mounting the spout apparatus on a pressure-driven sliding linear actuator (Festo) controlled by two solenoids (Parker). Licks were detected using an infrared emitter/receiver pair (Digikey) mounted on either side of the retractable lick spout. Rewards consisted of 5-8 μl water and punishments consisted of a white noise auditory stimulus alone (early training) or white noise plus 1-3 μl of 5mM quinine hydrochloride (Sigma) in water (late training). Behavioral training and testing was implemented with custom software

written in Matlab (Mathworks). Drifting grating stimuli were presented with the Psychophysics Toolbox (Brainard, 1997). Mice were trained during the light cycle. The stimulus consisted of sine wave gratings (spatial frequency: 0.05 cycles deg^{-1}; temporal frequency: 2 Hz) drifting at either 0 degrees (target) or 90 degrees (non-target) away from vertical. These stimuli were chosen to drive distinct groups of visual neurons with roughly equal strength. Thus, any large differences in stimulus selectivity observed in cortical neurons are not likely the result of stimulus strength.

Mice were trained in successive stages, with advancement to the next stage contingent on correct performance: 1) Mice received reward any time they licked the spout. 2) Trial structure was initiated by having an auditory cue tone, followed by a visual stimulus (100% targets), followed by an inter-trial interval. Mice were only rewarded for licks during the visual stimulus. 3) Once mice exhibited preferential licking during the stimulus compared to inter-trial interval, the target rate was reduced over several sessions from 100% to 50%. At this point, the non-target was a static grating orientated orthogonally to the target. Licks during non-targets were punished with white noise or white noise plus quinine. 4) Once mice exhibited the ability to discriminate target and non-target gratings ($d' > 1$ and $R_{HIT} - R_{FA} > 30\%$ for consecutive sessions, where R_{HIT} and R_{FA} are the hit and false alarm rate, respectively), the temporal frequency of the non-target grating was increased. 5) Spout withdrawal was introduced. At first the spout was extended within range before the stimulus appeared, then spout extend time was gradually delayed until after the stimulus had turned off (i.e., 0 s delay). 6) Finally, the variable delay period was gradually increased to 0/3/6 s. Mice that failed to fully learn the task within 150 sessions or showed signs of infection were removed from the study. In total, we removed 11 mice from the study before data collection was complete: 3 mice for failure to consistently lick the spout for reward (stage 1), 3 mice for failure to progress during the visual discrimination phase (stage 3), 1 mouse for failure to progress at the variable delay stage (stage 6), 1 mouse that showed signs of infection, and 4

mice that completed behavioral training but either had poor viral expression or cloudy windows after surgery.

Once mice reached high levels of performance at the final stage of the task (d' > 1.5 and R_{HIT} - R_{FA} > 50%), they were removed from water restriction for window implantation. After recovery from window implantation surgery, they were re-trained to a level of high performance (2-7 days) before beginning experimental sessions. Any sessions with poor performance were discarded (minimum performance criterion: d' > 1 and R_{HIT} - R_{FA} > 30%).

2.3.3 Two-photon imaging

GCaMP6s fluorescence was imaged 14-35 days after virus injection using Prairie Ultima IV 2-photon microscopy system with a resonant galvo scanning module (Bruker). For fluorescence excitation, we used a Ti-Sapphire laser (Mai-Tai eHP, Newport) with dispersion compensation (Deep See, Newport) tuned to λ = 910 nm. For collection, we used GaAsP photomultiplier tubes (Hamamatsu). To achieving a wide field of view, we used a 16x/0.8 NA microscope objective (Nikon), which was mounted on a Z-piezo (Bruker) for volume scanning. Resonant scanning was synchronized to z-piezo steps in the acquisition software for volume scanning. For volume scanning, four 441 x 512 pixel imaging planes separated by 20 or 25 μm were imaged sequentially at a stack rate of 5 Hz for 10 min imaging sessions. Occasionally, very bright neurons were visible across multiple planes. To exclude redundant sampling of the same neuron, the Pearson correlation coefficients of the fluorescence traces of all pairs of neurons within a recording were calculated. Neuron pairs with a correlation coefficient > 0.5, and an inter-ROI distance <12.5 μm in the XY plane were considered to be redundant and the ROI with the lower average fluorescence signal (more likely out of plane) was removed. Even with virus titer dilution, a small number of nucleus-filled neurons were observed in most experiments (Chen et al., 2013), but they comprised

a small percentage of neurons and generally did not exhibit significant task-driven responses (Harvey et al., 2012). Laser power ranged from 40-75 mW at the sample depending on GCaMP6s expression levels. Photobleaching was minimal (<1% min^{-1}) for all laser powers used. A custom stainless steel plate (eMachineShop.com) attached to a black curtain was mounted to the head plate before imaging to prevent light from the visual stimulus monitor from reaching the PMTs. During imaging experiments, the polypropylene tube supporting the mouse was suspended from the behavior platform with high tension springs (Small Parts) to dampen movement.

2.3.4 Image analysis

Images were acquired using PrairieView acquisition software and sorted into multi-page TIF files. All subsequent analyses were performed in MATLAB (Mathworks). First, images were corrected for X-Y movement by registration to a reference image (the pixel-wise mean of all frames) using 2-dimensional cross correlation. Movements in the Z-dimension were rare in normal imaging conditions, although movements were sometimes observed during licking. In order to prevent licking-related artifacts from being identified as significant responses, we adapted our response significance test to exclude short-duration changes in fluorescence (see Section 2.3.5).

To identify responsive neural somata, a pixel-wise activity map was calculated as previously described (Ahrens et al., 2012). Neuron cell bodies were identified using local adaptive threshold and iterative segmentation. Automatically-defined ROIs were then manually checked for proper segmentation in a MATLAB-based graphical user interface (allowing comparison to raw fluorescence and activity map images). To ensure that the response of individual neurons was not due to local neuropil on somatic signals, a corrected fluorescence measure was estimated according to previously described methods (Chen et al., 2013) as $F_{corrected\_soma}(t) = F_{raw\_soma}(t) - 0.7 \times F_{neuropil}(t)$, where $F_{neuropil}$ was the defined as the fluorescence in the region 0-15 μm from the ROI border (excluding other ROIs). The $\Delta F/F$ (corrected and uncorrected) for each neuron was

calculated as $\Delta F/F_t = (F_t - F_0)/ F_0$, with F_0 defined as the mode of the raw fluorescence density distribution.

2.3.5 Analysis of task-driven responses

After image preprocessing and $\Delta F/F$ extraction, traces were sorted into matrices by trial type (hit, miss, correct reject, false alarm). Testing neurons for significant responses was complicated by the large number of neurons (>9,000) and the number of samples per trial. Using a traditional threshold approach (with responses considered significant if they pass a threshold of multiple SDs from baseline) would yield either an unreasonable number of false positives (if a low threshold was used) or an unreasonable number of false rejections (if a high threshold was used). The GCaMP6s indicator has a long decay time constant ($\tau_{1/2}$ > 1 s), with fluorescence transients staying above baseline for >3 s seconds for even a single action potential (Chen et al., 2013). Thus, to capture weak but reliable trial-locked activity while excluding artifacts, we used a low significance threshold (sample different from baseline at $p < 0.05$, Wilcoxon signed-rank test) but required at least 10 samples to be significant in the same direction for at least 2 (of 3) delays for either hit or CR trials to be considered statistically significant (probability of significant response for any one neuron purely by chance $<10^{-9}$). Since reliable GCaMP6s signals will exhibit slow decay over 10 or more samples, this approach will allow genuine calcium signals to emerge as significant, while chance fluctuations in fluorescence and short-duration movement artifacts (such as during licking) will not pass the significance threshold. All neurons that were significantly different from baseline according to these criteria in either the positive (enhanced activity) or negative (suppressed activity) were included in further analysis. To ensure the neuropil correction procedure did not create artificial responses, we required that both uncorrected $\Delta F/F$ and corrected $\Delta F/F$ exhibit significant changes for neurons to be included. Once neurons were determined to have significant responses, the corrected $\Delta F/F$ was used for all further analyses.

For some analyses, the $\Delta F/F$ responses were normalized by subtracting the baseline response (1 s before stimulus onset) and dividing by the maximum (for enhanced neurons) or minimum (for suppressed neurons) trial-averaged $\Delta F/F$ for the neuron.

To investigate the clustering of task-driven responses (**Figure 2.3**), we first de-noised and reduced the dimensionality of the data by taking the first 20 principal components (explaining 98% of the variance). To cluster the data, we computed linkages using Ward's method (with Euclidean distance). The dendrogram revealed distinct groups with high inter-cluster distance (**Figure 2.3A**). Averaging the responses of neurons within these clusters yielded distinct response types that corresponded well to the range of observed responses (**Figure 2.3B)**. To determine the number of clusters (K = 6), we iteratively increased the number of clusters until the average traces began showing highly overlapping responses (at K > 6). Manual inspection revealed a small number of neurons within each cluster that appeared to be miscategorized. To improve categorization (**Figure 2.3C)**, we computed the Pearson correlation coefficient between the response of each neuron and the mean of each of the six response types. Neurons were re-categorized based on the response type with highest correlation to their average response, yielding distinct clusters with high face validity and internal consistency.

The trial selectivity index (**Figure 2.4C**) was computed as $SI = (R_{pref} - R_{non\text{-}pref})/(R_{pref} + R_{non\text{-}pref})$, where R_{pref} and $R_{non\text{-}pref}$ are the responses to the preferred and non-preferred trials, respectively. The modal latency of the population (**Figure 2.4D, Figure 2.5D**) was estimated by first taking the DF/F of each neuron, and upsampling (via linear interpolation) the mean $\Delta F/F$ signal to 1000 Hz. We then determined the point at which the mean $\Delta F/F$ signal for each neuron went above or below baseline in a sustained manner ($>1\sigma$ from mean baseline $\Delta F/F$ for 1000 consecutive samples, or 1 s in trial time) and used the first sample as the estimated latency for that neuron. Since the latency distribution of the population was highly skewed in higher regions, we used the first mode of the latency density distribution to describe the onset of population

activity (full latency distributions for 6 s delay trials shown in **Figure 2.5D**). Note that these latencies are computed on upsampled data based on a linear interpolation, so the values should be considered approximate. To measure whether the neurons within a region exhibited delay modulation (**Figure 2.4E**) we fit the slope of the integrated $\Delta F/F$ across increasing delay durations using a first order polynomial. We then calculated the delay modulation index (DMI) as DMI = $(n_{pos}-n_{neg})/(n_{pos}+n_{neg})$ for enhanced populations and DMI = $(n_{neg}-n_{pos})/(n_{pos}+n_{neg})$ for suppressed populations, where n_{pos} is the number of neurons with a positive slope and n_{neg} is the number of neurons with a negative slope.

2.3.6 Population encoding of task-related variables

In order to determine how well we could decode the stimulus identity for populations of neurons (**Figure 2.6B**), we measured the discriminability of population responses on error trials from correct trials with a different stimulus but the same response. For example, to compare miss and correct reject (CR) population responses (different stimuli, same response), for each time point we first calculated the Euclidean distances between the neural population response vector at each CR trial *t* and a template CR response vector (average of all CR trials except *t*) and calculated the Euclidean distances between each miss trial *t* and the template CR response (average of all trials). We then generated a receiver operating characteristic curve from the distances to determine the discriminability of the miss from CR responses. By taking the area under the receiver operating characteristic (auROC; (Britten et al., 1992) for each time point, we could quantify the performance of an ideal observer in discriminating stimulus identity based purely on the population responses as a function of time during the trial. For behavioral choice encoding (**Figure 2.6C**), the same analysis was carried out comparing error-correct trial pairs with the same stimulus but different response. To reduce noise from experiments with very few error trials, only experiments that had at least three error trials of each type (miss, false alarm) were

included in the analysis. To determine how coding varied as a function of neural population size, we averaged the performance across 1000 iterations using populations of 1, 5, 10, 50, or 100 randomly selected neurons (**Figure 2.6D**). For significance testing, we calculated the auROC for 1000 permutations with shuffled trial labels. Individual time points were considered to predict stimulus/choice if the mean auROC was greater than the 95% confidence interval of the shuffled permutations for at least 3 consecutive time points (to compensate for multiple comparisons).

2.3.7 General statistics

Data groups were tested for normality using the Kolmogorov-Smirnov test and then compared with the appropriate tests (t-tests, Wilcoxon rank-sum or Wilcoxon signed-rank tests, all two sided). Bonferroni correction was used for multiple comparisons. Bootstrap estimates of s.e.m. were calculated as the standard deviation of values evaluated in 1,000 bootstrap iterations, obtained by randomly re-sampling with replacement from the original values. Due to very large sample sizes, very small p-values ($<10^{-9}$) were approximated as $p < 10^{-9}$ as a lower bound on reasonable probabilities. Sample sizes were not explicitly estimated, as even a single session generally had sufficient samples for the statistical tests used. However, to ensure that results were replicable between sessions and mice, we included 4-6 mice for each region in both imaging and inactivation experiments.

2.4 Results

2.4.1 Calcium imaging during a memory-guided sensorimotor task

We trained head-fixed mice to perform a visual discrimination task with a memory-guided response (**Figure 2.1A**). In this task, water-restricted mice were presented a 2 s drifting grating

stimulus at one of two orientations (0°, 90° from vertical), followed by a variable delay period (0-, 3-, or 6-s), at which point a lick spout was moved rapidly into reach with a linear actuator for 1.5 s (**Figure 2.1B**, bottom). Licking on "go" trials (horizontal grating drifting toward 0° from vertical) was rewarded with 5-8 μl water (hit), while licking on "no-go" trials (vertical grating drifting toward 90° from vertical) was punished with 2 μl water containing 5mM quinine hydrochloride (false alarm; **Figure 2.1B**, top). This structure allowed the separation of each trial into "stimulus", "delay", and "response" epochs. Notice that the stimulus-choice association is fixed, so correct performance does not require memory of the stimulus during the delay period (memory of the planned response is sufficient). After extensive training (117 ± 11 behavioral sessions, 299 ± 25 trials per session; mean ± s.e.m.), mice reliably exhibited strong differences in licking between go and no-go trials (**Figure 2.1C**, top) for all delay period durations. We applied an *a priori* exclusion criteria that any mice licking continuously throughout the delay period on target trials would be excluded from the study, since this strategy would possibly obviate the short-term memory component of the task. Mice did exhibit a bias toward licking (as observed previously with go/no-go tasks; (Huber et al., 2012; O'Connor et al., 2010), so we quantified their performance using d-prime rather than percent correct to account for motivation and criterion (Carandini and Churchland, 2013) (**Figure 2.1C**, bottom). Performance decreased slightly with longer delays, but was well above chance for all delay durations (0 s Delay, $p < 10^{-9}$; 3 s Delay, $p < 10^{-9}$; 6 s Delay, $p < 10^{-9}$; *t*-test, $n = 8$ mice across 80 sessions).

We focused our experiments on three cortical regions we expected to be important for performance of this task: (1) the primary visual cortex (V1), which is known to be important for orientation discrimination (Glickfeld et al., 2013); (2) the posterior parietal cortex (PPC), which receives extensive input from visual regions (Harvey et al., 2012; Oh et al., 2014; Pho et al., 2015), projects to motor regions (Wang et al., 2012), and has been implicated in sensorimotor decision tasks (Andersen and Cui, 2009; Gold and Shadlen, 2007; Hanks et al., 2015;

McNaughton et al., 1994; Nitz, 2006; Raposo et al., 2014; Shadlen and Newsome, 2001; Whitlock et al., 2008) and (3) the frontal motor cortices (fMC), which include regions known to be crucial for voluntary licking behaviors (Guo et al., 2014b; Komiyama et al., 2010).

We identified parietal cortex on the basis of stereotaxic coordinates from previous studies (Harvey et al., 2012). Note that this region has weak visual responses and has also been classified as a secondary visual region (AM) (Garrett et al., 2014; Wang et al., 2012). However, in addition to visual inputs, retrograde tracing from our lab (unpublished results) and others (Harvey et al., 2012) has revealed that the region receives input from auditory, somatosensory, secondary motor, and frontal cortices, as well as the lateral posterior thalamic nucleus. Since the parcellation of rodent frontal and motor cortices is a subject of debate in the field (Brecht, 2011), and both medial and lateral regions have been implicated in licking (Komiyama et al., 2010) we define fMC on the basis of stereotaxic coordinates (including primary and secondary motor regions) and remain agnostic as to its precise homology with primate frontal and motor cortices. Nonetheless, several studies have indicated that rodent fMC plays an important role in tasks involving perceptual decisions and memory (Erlich et al., 2015; Guo et al., 2014b; Kepecs et al., 2008; Li et al., 2015).

To measure neural activity in layer 2/3 of cortex during task performance, a craniotomy was made over one or more of regions V1, PPC and fMC (**Figure 2.1D**) after completion of training. Stereotaxically-guided microinjections of adeno-associated virus (AAV) containing the genetically-encoded calcium indicator GCaMP6s (Chen et al., 2013) were made in V1 (n = 5 mice), PPC (n = 6) and/or fMC (n = 4), and sealed cranial windows were made over V1 and PPC (**Figure 2.1E**) or fMC. To increase the number of recorded neurons, we used a volume scanning approach (Kampa et al., 2011) that allowed us to image four imaging planes (separated by 20-25 μm) to simultaneously sample hundreds of GCaMP6s-infected neurons within an 850 μm x 850 μm x 60-75 μm volume at a sample rate of 5 Hz (**Figure 2.1F**). Images were corrected for X-Y movement and regions of interest were assigned to cell somata based on a pixel-wise activity

map calculation, yielding fluorescence traces from individual neurons that were active during the imaging session (26 sessions, active neurons per session: 352 ± 40, mean ± s.e.m.). A total of 9,150 active neurons were imaged in regions V1 (*n* = 2,695 neurons), PPC (*n* = 3,552), and fMC (*n* = 2,903), of which 3,049 (33.3%) exhibited trial-locked responses significantly different from baseline (see **Experimental Procedures** for inclusion criteria; **Figure 2.1G**). Pilot experiments in transgenic mice expressing tdTomato in PV+ and SOM+ interneurons revealed that interneuron calcium signals measured with volume scanning were considerably smaller than tdTomato-negative (putatively excitatory) neurons (though see (Peron et al., 2015), and therefore the vast majority of task-responsive cells were likely excitatory pyramidal neurons.

2.4.2 Single neurons exhibit heterogeneous trial-evoked responses

To investigate how regions V1, PPC, and fMC encode task-relevant variables, we analyzed the activity of single neurons in each region during task performance (**Figure 2.2**). The majority (63%, **Figure 2.3,** see below for description of classification procedure) of V1 neurons with increased task-evoked activity exhibited robust and reliable responses during the stimulus epoch but not during other epochs of the task, with similar responses across 0, 3, and 6 s delays (**Figure 2.2A,B**). However, there was also a sizable fraction (37%) of neurons that exhibited activity during the other task epochs, particularly during the response epoch (**Figure 2.2C**). Interestingly, there was an additional fraction (53% of responsive neurons) of neurons that were suppressed throughout the stimulus and delay periods in a delay duration-dependent fashion (**Figure 2.2D**). To our knowledge, neurons exhibiting delay-dependent suppression have not been previously described in V1 despite their prevalence in our sample.

Neurons in PPC were even more heterogeneous, with a large number (48% of enhanced neurons) of neurons exhibiting activity during both stimulus and response epochs (**Figure 2.2E,F**), as well as some neurons (11% of enhanced neurons) exhibiting delay-dependent enhanced

activity (**Figure 2.2G**), as has previously been observed in primate experiments in delayed-response tasks (Chafee and Goldman-Rakic, 1998; Constantinidis and Steinmetz, 1996). Finally, we again saw neurons that exhibited suppressed delay-dependent activity during the task (**Figure 2.2H**), though they were much less prevalent compared to V1 (9% vs. 53% for PPC and V1, respectively).

Neurons in fMC likewise exhibited heterogeneous responses during multiple task epochs. As expected in motor regions, a plurality (47% of enhanced neurons) of neurons responded solely during the response epoch (**Figure 2.2I**). However, a substantial fraction (33% of enhanced neurons) of neurons responded not only during the response epoch, but also in a sustained manner throughout the delay period (**Figure 2.2J**), including neurons that exhibited sustained activity during the delay period of the task, and actually showed decreased activity during the response period (**Figure 2.2K**). Note that the sustained activity cannot be attributed to the slow decay of the GCaMP6 indicator after the initial stimulus response, as the duration of the sustained activity varies parametrically with the duration of the delay period. Sustained activity was also observed using a calcium indicator with a faster time course (GCaMP6f, data not shown). Finally, we again observed a subset (39% of responsive neurons) of neurons with suppressed delay-dependent activity in fMC (**Figure 2.2L**).

To characterize the diversity of trial-evoked responses and their relative prevalence in each region, we clustered the neurons by their response characteristics. We pooled the significant responses (enhanced and suppressed) from all three areas, and then used PCA-based de-noising and hierarchical clustering to delineate distinct response types (see Materials and Methods). Hierarchical clustering revealed six distinct response types, four of them exhibiting enhanced activity and two exhibiting suppressed activity (**Figure 2.3A**). To reveal the archetypical response of each group, we averaged the normalized responses of all neurons within each cluster, revealing four enhanced response classes: stimulus-driven (*Stim only*), response-driven

(*Resp only*), stimulus- and response-driven (*Stim+Resp*), and delay-driven (*Delay*), as well as two suppressed response classes: delay-sensitive stimulus-driven (*Supp. Delay early*) and delay-sensitive response-driven (*Supp. Delay late*; **Figure 2.3B**). These clusters have strong face validity, as the average responses (**Figure 2.3B**) are qualitatively similar to observed single-neuron responses (**Figure 2.2**). Although neurons that were active during the response epoch were predominantly selective for 'go' trials, there were response-driven neurons that were selective for 'no-go' trials as well (data not shown), suggesting that the response-driven activity is not entirely related to motor activity. There were clear differences in both enhanced and suppressed response types between regions, such as V1 containing a large number of '*Stim only*' and 'Supp. *Delay early*' neurons, while fMC had greater numbers of '*Delay*', '*Resp only*', and '*Supp. Delay late*' neurons (**Figure 2.3C**). However, there was a surprising amount of heterogeneity as well, with all three regions containing a fraction of almost all response types.

2.4.3 Suppressed neurons exhibit non-selective responses

We next investigated what role neurons exhibiting suppressed activity (975/3049 neurons, 32.0%; **Figure 2.2D,H,L**) play in encoding task-relevant information. We found that the enhanced and suppressed neurons across regions exhibited striking differences in selectivity. The vast majority of enhanced neurons showed a strong preference for a particular trial type (hit vs. correct reject, only correct trials included in analysis), with little or no response to the non-preferred trial type (**Figure 2.4A**). However, the majority of suppressed neurons showed very similar responses to both trial types, with little difference between "preferred" and "non-preferred" trial types (**Figure 2.4B**). We used a selectivity index ranging from 0 (responds equally to preferred and non-preferred stimuli) to 1 (only responds to preferred stimuli) to quantify the selectivity of each neuron (see **Materials and Methods**), revealing that enhanced neurons were significantly more selective than suppressed neurons in all three regions (**Figure 2.4C**; V1: Enh. median SI = 0.67, Supp.

median SI = 0.16, $p < 10^{-9}$; PPC: Enh. median SI = 0.91, Supp. median SI = 0.31, $p < 10^{-9}$; fMC: Enh. median SI = 0.99, Supp. median SI = 0.53, $p < 10^{-9}$; Wilcoxon rank-sum test).

One possibility for the presence of the suppressed responses is that local inhibition is being recruited by the increased activity of the enhanced neurons during task performance. If this were the case, we would expect that the time course of the suppressed responses would closely follow that of the enhanced responses. We estimated the response latency of the suppressed neurons, and found that the latencies were slow (>400 ms) and uncorrelated to the enhanced population latency in the cortical region (**Figure 2.4D**). We also found that suppressed neurons exhibited suppression throughout the delay period (in a duration-dependent manner; **Figure 2.2D,H,L**), even when this pattern was not present in the enhanced neurons of the same region (**Figure 2.4E**). Taken together, these findings suggest that the suppressed responses are not simply reflecting inhibition from local excitatory responses, but rather are the result of more complex dynamics; possibly low-latency, delay-dependent inputs from distal regions.

2.4.4 Enhanced neurons show regional differences despite local heterogeneity

The lack of trial type selectivity observed in suppressed neurons indicates that there will be little information about stimulus identity or future motor response in the activity of these neurons. However, the enhanced neurons are highly selective, and likely represent these task-relevant variables in a more robust manner. To further investigate the role of enhanced neurons in different cortical regions, we investigated the population responses within regions by pooling the neurons across imaging sessions from different animals.

To investigate the encoding of task variables at the population level, we averaged and normalized the preferred responses of enhanced task-driven neurons across all correct trials of equivalent delay (hit or correct reject, depending on preferred response) and sorted them first by cortical region, and then by latency to peak response (**Figure 2.5A**). While V1 neurons preferred

go and no-go trials in roughly equal numbers, we found that PPC and fMC neurons were highly biased toward the go trials (**Figure 2.5B**). Since the target and non-target stimuli have similar visual saliency, this suggests that these regions may encode task-related variables other than stimulus identity. Indeed, although the responses within regions were heterogeneous (**Figure 2.3**), there were clear differences in the average responses between regions (**Figure 2.5C**). Specifically, V1 neurons were predominantly active during the stimulus epoch, while PPC and fMC exhibited more heterogeneous responses spanning stimulus, delay, and response epochs. Although sustained delay-period activity could be observed across all cortical regions, it was most prevalent in fMC, both in single-unit (**Figure 2.3C)** and population **(Figure 2.5C)** activity. Finally, there were considerable differences in latency to significant response, with the first mode of the latency distribution increasing from V1 to PPC to fMC for all delay durations (V1: 89 ± 10ms, PPC: 244 ± 15 ms, fMC: 827 ± 177 ms, mean across delays; regions all significantly different, V1 x PPC: $p < 10^{-9}$, V1 x fMC: $p < 10^{-9}$, PPC x fMC: $p < 10^{-9}$, Wilcoxon rank-sum test; **Figure 2.5D**).

2.4.5 Population error-trial analysis reveals distinct encoding dynamics in each region

To further investigate how neural activity might reflect the encoding of the stimulus identity and behavioral choice (including planning of the response before it is made), we analyzed the modulation of responses in all regions during error trials. We hypothesized that if neurons were simply encoding stimulus identity, there would be no difference in activity during correct trials and error trials containing the same visual stimulus (i.e., hit vs. miss; correct reject vs. false alarm). Indeed, V1 neurons exhibit similar responses to correct and error trials with the same stimulus, while neurons in PPC and fMC show very different responses on the two trial types (**Figure 2.6A**, left and middle panels corresponding to miss vs. hit trials; Pearson correlation between miss and hit trials, V1: 0.47 ± 0.03, PPC: 0.05 ± 0.01, fMC: 0.06 ± 0.02). Similarly, we hypothesized that if neurons encoded the animal's behavioral choice independently of the stimulus shown, we would

expect there to be little difference between correct trials and error trials with the same motor response (i.e., hit vs. false alarm; correct reject vs. miss). Neuronal responses in PPC and fMC appear to exhibit more similarity than V1 for trial types with the same motor response (**Figure 2.6A**, middle and right panels corresponding to hit vs. false alarm trials; Pearson correlation between hit and false alarm trials, V1: 0.23 ± 0.03, PPC: 0.47 ± 0.01, fMC: 0.46 ± 0.02).

To quantify the encoding of task variables as a function of time in each region, we used an ideal observer analysis to determine how well we could decode the stimulus identity independent of the behavioral choice (**Figure 2.6B**) and behavioral choice independent of stimulus identity (**Figure 2.6C**) using only the responses from randomly selected populations of neurons (ranging from population size of 1 to 100, see **Materials and Methods** for details of decoding procedure). In V1, stimulus identity could be decoded with perfect accuracy given a sufficient number of neurons (**Figure 2.6B**, left), while behavioral choice was only weakly encoded (**Figure 2.6C**, left). The stimulus identity encoding peaked during the visual stimulus and then gradually declined throughout the trial (statistically significant from baseline from 0.2 - 9.2 s, mean value greater than 95% CI of shuffled permutations for consecutive time points, see **Materials and Methods**), likely due to the slow decay of the GCaMP6s indicator, while the behavioral choice encoding gradually increased, peaking during the response period (statistically significant from 4.8 – 6.8 s and 8.4 - 10.0 s; **Figure 2.6D**, left). Note that in all three regions, the increase in stimulus identity encoding during the response epoch is likely the byproduct of reward or punishment (which is directly associated with the stimulus identity) influencing motor activity. In PPC, both stimulus identity and behavioral choice could be decoded with moderate success (**Figure 2.6B,C**, middle), although the dynamics were noticeably different. Specifically, while the stimulus identity was predominantly encoded early in the trial (statistically significant from 0.6 – 3.4 s and 7.2 - 10.0 s), the behavioral choice encoding slightly lagged the stimulus identity encoding, and remained high throughout the trial, peaking during the early part of the licking

response (statistically significant from 1.4 - 10.0 s; **Figure 2.6D**, middle). Note that the stimulus encoding is much weaker in PPC than in V1, likely due to the large number of neurons that jointly-encode stimulus and choice information (Park et al., 2014; Raposo et al., 2014); Pho et al., 2015). Finally, in fMC the stimulus identity could not be decoded above chance for the majority of the trial (**Figure 2.6B**, right), while the behavioral choice could be decoded with greater accuracy in fMC than in other regions (**Figure 2.6C**, right). Thus, while stimulus identity decoding was only above chance during the response epoch (statistically significant from 8.8 - 9.4 s, see caveat on stimulus identity encoding during response epoch above), the behavioral choice could be decoded throughout the delay and response epochs (statistically significant from 3.0 - 10.0 s; **Figure 2.6D**, right).

Taken together, these results suggest a working hypothesis of the roles that neural activity in regions V1, PPC, and fMC play in the performance of the task. Specifically, the results suggest that sensory input is primarily processed during the stimulus epoch, first in V1, then subsequently in PPC. Neural activity related to the behavioral choice arises in PPC and fMC shortly after the peak in stimulus identity coding, and is sustained in both regions throughout the delay and response epochs. This implies that stimulus identity is rapidly transformed into a behavioral choice within the stimulus epoch (possibly within PPC), and then the behavioral choice is maintained in higher regions (potentially in both PPC and fMC) until the relevant motor action is performed (**Figure 2.7**).

2.5 Discussion

Understanding the circuits underlying the conversion of sensory input into motor action is a major goal of neuroscience research (Andersen and Cui, 2009; Gold and Shadlen, 2007; Romo and de Lafuente, 2013). By measuring the activity of neurons in multiple cortical regions during a memory-guided sensorimotor decision task, this study has made the following contributions.

First, large-scale 2-photon calcium imaging revealed that despite differences in single-unit activity between regions, the responses within a region were surprisingly heterogeneous, with a representation of all classified response types in all regions (**Figure 2.3**). For example, in V1 we observed many neurons that were active during the response epoch in addition to the expected neurons active during the stimulus period, and vice versa in the frontal motor cortices. This study also shows, for the first time in mice, that a fraction of neurons (mostly in PPC and fMC) exhibit sustained activity that varies parametrically with delay duration, in line with earlier work from non-human primates (Funahashi et al., 1989; Fuster and Alexander, 1971; Kojima and Goldman-Rakic, 1982).

One difference between our study and many previous delayed-response studies was the use of a go/no-go design rather than an alternative choice design. Although we used this design primarily for practical reasons (to reduce training time), we found interesting biases in the cortical responses resulting from the response asymmetry (**Figure 2.5B**). In particular, while V1 exhibited a moderate bias toward Hit trials (65.3% target-preferring), PPC and fMC exhibited much stronger biases toward Hit trials (95.5% and 97.2% target-selective, respectively), although there was a small fraction of neurons that showed robust responses selective for CR. Given that the false alarm rate exceeded the miss rate at all delays, it is surprising that there is not more activity reflecting the suppression of licking on non-target trials. In future studies, it would be interesting to probe other regions involved in motor and cognitive control, such as striatum and prefrontal cortex, to see if there are signals corresponding to the suppression of licking during correct reject trials.

The sustained delay-period activity we observed contrasts with several recent studies in which neurons recorded in parietal and prefrontal cortices exhibited sequential neural activity (Baeg et al., 2003; Fujisawa et al., 2008; Harvey et al., 2012). It has been unclear whether the different modes of neural activity (steady-state vs. sequential) observed during short-term

memory depend on the behavioral task, or whether the sequential activity had simply been obscured by lower-throughput methods of collecting and analyzing the data. Using volumetric calcium imaging to sample large populations of neurons, we found evidence in our task for steady-state representation of motor choice, with the vast majority of delay-sensitive neurons becoming active early in the trial and exhibiting persistent or ramping activity throughout the delay period (**Figure 2.2G, J, K; Figure 2.3**). We suspect that the nature of the mnemonic information in the behavioral task has a strong bearing on the mode of activity observed. In tasks where neurons exhibit steady-state activity, such as our study and others (Funahashi et al., 1989; Fuster and Alexander, 1971; Kojima and Goldman-Rakic, 1982; Shadlen and Newsome, 2001), the encoded motor plan is held constant throughout the delay period (e.g., lick when the spout comes forward; make a saccade to a fixed location). In contrast, in tasks where neurons exhibit sequential activity (Baeg et al., 2003; Fujisawa et al., 2008; Harvey et al., 2012) the encoded motor plan may vary as a function of the trial as the animal navigates through real or virtual space. Both types of memory are important for understanding complex behavior, and warrant further study of the underlying neural architecture.

A second contribution is the finding that a sizable fraction of the significantly responding neurons (32.0% of all significantly-responding neurons) exhibited suppressed rather than enhanced responses during task performance (**Figure 2.2**, **Figure 2.3**). However, these neurons were generally much less selective than enhanced neurons (particularly in V1 and PPC), and as a result appear to play little role in encoding task-relevant variables (**Figure 2.4**). Without a method for specifically manipulating the activity of the suppressed neurons, it is difficult to unambiguously discern what role they may play in task performance. One possibility is that suppression of spontaneously active neurons may reduce any ongoing task-irrelevant activity, and thereby improve the readout of task-relevant population activity by higher-order neurons.

Third, we found that despite the heterogeneity observed in the single-unit activity, when the populations of neurons within each region are analyzed as a whole, they exhibit distinct coding dynamics of task-relevant variables (**Figure 2.6**). Specifically, the V1 population predominantly encodes the stimulus identity, particularly early in the trial. Similar to previous research in non-human primates (Park et al., 2014) and rodents (Raposo et al., 2014), we find multiplexed encoding of stimulus identity and behavioral choice in PPC. However, we find that the dynamics of the stimulus identity and behavioral choice encoding in PPC are distinct, with choice encoding lagging slightly, and lasting throughout the delay period. The fMC predominantly encodes behavioral choice, beginning early in the trial and lasting throughout the delay period until after the motor response is made, in accord with choice-related activity observed in premotor regions in earlier rodent (Erlich et al., 2011; Guo et al., 2014b; Li et al., 2015) and non-human primate studies (Hernandez et al., 2010).

A few caveats are important to note in the interpretation of this study. One persistent issue in delayed-response tasks (in which the correct response is evident early in the trial) is that the subject could adapt a behavioral strategy in which they make a micropostural preparation (e.g., movement of the eyes, changes in body orientation or tongue position) to stand in for the eventual motor action. This in turn may preclude the necessity for actively maintaining a motor plan throughout the delay period. This problem has been recognized for some time (Kojima and Goldman-Rakic, 1982), and we attempted to address the possibility using video analysis (see **Supplementary Materials**), but it remains challenging to control for the wide range of potential covert movements. Disentangling the effects of covert preparations on premotor and motor activity will be an important step in understanding the neural basis of motor planning. Use of more advanced behavioral paradigms, such as a delayed match-to-sample task, can additionally aid in isolating periods of sensory memory from motor planning.

A second caveat is that the asymmetry of the go/no-go design can muddle interpretation of stimulus and choice information, as assessed via comparison with error trials (**Figure 2.6**). Errors in a go/no-go paradigm can result from changes in arousal, motivation, or impulsivity that alter the willingness of the animal to respond, and therefore our measurement of "choice"-related information may also reflect these non-specific factors. While go/no-go designs have been commonly used for rodent studies (Huber et al., 2012; Lee et al., 2012; Pinto and Dan, 2015; Pinto et al., 2013; Sippy et al., 2015), forced-choice tasks are also amenable to rodents (Brunton et al., 2013; Guo et al., 2014a; Raposo et al., 2014) and have the advantage of being immune to these issues of motivation and criterion (Carandini and Churchland, 2013).

Finally, this study focused on cortical regions located on the dorsal surface of the mouse brain, as they are readily accessible to both 2-photon calcium imaging and optogenetic inhibition. However, there are likely to be additional cortical and subcortical regions involved in this task, including prefrontal cortex (Funahashi et al., 1989; Fuster and Alexander, 1971; Kojima and Goldman-Rakic, 1982), thalamus (Fuster and Alexander, 1973), superior colliculus (Kopec et al., 2015), and basal ganglia (Ding and Gold, 2012; Kawagoe et al., 1998). These regions could be further investigated with a similar approach as invasive (Andermann et al., 2013; Barretto and Schnitzer, 2012; Zorzos et al., 2012) and noninvasive (Filonov et al., 2011; Mittmann et al., 2011; Prakash et al., 2012) techniques for optical interrogation of deeper structures become available.

2.6 Contributions and Acknowledgements

M.J.G. designed experiments with input from G.N.P. and M.S.; M.J.G. performed the imaging experiments; M.J.G., G.N.P., and JW performed the behavioral training; G.N.P. developed the behavior software; M.J.G. analyzed the data with comments from G.N.P. and M.S.; M.J.G., G.N.P., and M.S. wrote the manuscript.

We thank J. Sharma, S. Ramirez, B. Crawford, A. Boesch, V. Li, C. Le, and T. Emery for technical assistance; S. El-Boustani and R. Huda for comments on the manuscript; L. L. Looger, J. Akerboom, D. S. Kim, and the Genetically-Encoded Neuronal Indicator and Effector (GENIE) Project at Janelia Farm Research Campus Howard Hughes Medical Institute for generating and characterizing GCaMP6 variants. This work was supported by the NIH (M.J.G., F32-EY023523 and K99-MH104259; M.S., R01-EY007023 and U01-NS090473); the NSF (G.N.P., Graduate Research Fellowship; M.S., EF1451125); the Simons Center for the Social Brain (M.S.), and the Picower Institute Innovation Fund (M.J.G.; M.S.).

2.7 References

Ahrens, M.B., Li, J.M., Orger, M.B., Robson, D.N., Schier, A.F., Engert, F., and Portugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish. Nature *485*, 471-477.

Andermann, M.L., Gilfoy, N.B., Goldey, G.J., Sachdev, R.N., Wolfel, M., McCormick, D.A., Reid, R.C., and Levene, M.J. (2013). Chronic cellular imaging of entire cortical columns in awake mice using microprisms. Neuron *80*, 900-913.

Andermann, M.L., Kerlin, A.M., and Reid, R.C. (2010). Chronic cellular imaging of mouse visual cortex during operant behavior and passive viewing. Frontiers in cellular neuroscience *4*, 3.

Andersen, R.A., and Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits. Neuron *63*, 568-583.

Baeg, E.H., Kim, Y.B., Huh, K., Mook-Jung, I., Kim, H.T., and Jung, M.W. (2003). Dynamics of population code for working memory in the prefrontal cortex. Neuron *40*, 177-188.

Barretto, R.P., and Schnitzer, M.J. (2012). In vivo optical microendoscopy for imaging cells lying deep within live tissue. Cold Spring Harb Protoc *2012*, 1029-1034.

Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. J Neurosci *31*, 913-921.

Brainard, D.H. (1997). The Psychophysics Toolbox. Spat Vis *10*, 433-436.

Brecht, M. (2011). Movement, confusion, and orienting in frontal cortices. Neuron *72*, 193-196.

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci *12*, 4745-4765.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95-98.

Carandini, M., and Churchland, A.K. (2013). Probing perceptual decisions in rodents. Nat Neurosci *16*, 824-831.

Chafee, M.V., and Goldman-Rakic, P.S. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. J Neurophysiol *79*, 2919-2940.

Chafee, M.V., and Goldman-Rakic, P.S. (2000). Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. J Neurophysiol *83*, 1550-1566.

Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V.*, et al.* (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295-300.

Compte, A., Brunel, N., Goldman-Rakic, P.S., and Wang, X.J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb Cortex *10*, 910-923.

Constantinidis, C., and Steinmetz, M.A. (1996). Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. J Neurophysiol *76*, 1352-1355.

de Lafuente, V., and Romo, R. (2005). Neuronal correlates of subjective sensory experience. Nature neuroscience *8*, 1698-1703.

di Pellegrino, G., and Wise, S.P. (1993). Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. J Neurosci *13*, 1227-1243.

Ding, L., and Gold, J.I. (2012). Separate, causal roles of the caudate in saccadic choice and execution in a perceptual decision task. Neuron *75*, 865-874.

Erlich, Jeffrey C., Bialek, M., and Brody, Carlos D. (2011). A Cortical Substrate for Memory-Guided Orienting in the Rat. Neuron *72*, 330-343.

Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. Elife *4*.

Filonov, G.S., Piatkevich, K.D., Ting, L.M., Zhang, J., Kim, K., and Verkhusha, V.V. (2011). Bright and stable near-infrared fluorescent protein for in vivo imaging. Nat Biotechnol *29*, 757-761.

Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. Nature *443*, 85-88.

Fujisawa, S., Amarasingham, A., Harrison, M.T., and Buzsaki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. Nat Neurosci *11*, 823-833.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J Neurophysiol *61*, 331-349.

Fuster, J.M., and Alexander, G.E. (1971). Neuron activity related to short-term memory. Science *173*, 652-654.

Fuster, J.M., and Alexander, G.E. (1973). Firing changes in cells of the nucleus medialis dorsalis associated with delayed response behavior. Brain Res *61*, 79-91.

Garrett, M.E., Nauhaus, I., Marshel, J.H., and Callaway, E.M. (2014). Topography and areal organization of mouse visual cortex. J Neurosci *34*, 12587-12600.

Glickfeld, L.L., Histed, M.H., and Maunsell, J.H. (2013). Mouse primary visual cortex is used to detect both orientation and contrast changes. J Neurosci *33*, 19416-19422.

Goard, M.J., Pho, G.N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife *5*, 558.523.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annual review of neuroscience *30*, 535-574.

Goldman, M.S., Levine, J.H., Major, G., Tank, D.W., and Seung, H.S. (2003). Robust persistent neural activity in a model integrator with multiple hysteretic dendrites per neuron. Cereb Cortex *13*, 1185-1195.

Guo, Z.V., Hires, S.A., Li, N., O'Connor, D.H., Komiyama, T., Ophir, E., Huber, D., Bonardi, C., Morandell, K., Gutnisky, D.*, et al.* (2014a). Procedures for behavioral experiments in head-fixed mice. PloS one *9*, e88678.

Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014b). Flow of cortical activity underlying a tactile decision in mice. Neuron *81*, 179-194.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220-223.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62-68.

Hernandez, A., Nacher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vazquez, Y., Camarillo, L., and Romo, R. (2010). Decoding a perceptual decision process across cortex. Neuron *66*, 300-314.

Huber, D., Gutnisky, D.A., Peron, S., O'Connor, D.H., Wiegert, J.S., Tian, L., Oertner, T.G., Looger, L.L., and Svoboda, K. (2012). Multiple dynamic representations in the motor cortex during sensorimotor learning. Nature *484*, 473-478.

Kampa, B.M., Roth, M.M., Gobel, W., and Helmchen, F. (2011). Representation of visual scenes by local neuronal populations in layer 2/3 of mouse visual cortex. Front Neural Circuits *5*, 18.

Kawagoe, R., Takikawa, Y., and Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. Nat Neurosci *1*, 411-416.

Kepecs, A., Uchida, N., Zariwala, H.A., and Mainen, Z.F. (2008). Neural correlates, computation and behavioural impact of decision confidence. Nature *455*, 227-231.

Kojima, S., and Goldman-Rakic, P.S. (1982). Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. Brain Res *248*, 43-49.

Komiyama, T., Sato, T.R., O'Connor, D.H., Zhang, Y.X., Huber, D., Hooks, B.M., Gabitto, M., and Svoboda, K. (2010). Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. Nature *464*, 1182-1186.

Kopec, C.D., Erlich, J.C., Brunton, B.W., Deisseroth, K., and Brody, C.D. (2015). Cortical and Subcortical Contributions to Short-Term Memory for Orienting Movements. Neuron *88*, 367-377.

Lee, S.H., Kwan, A.C., Zhang, S., Phoumthipphavong, V., Flannery, J.G., Masmanidis, S.C., Taniguchi, H., Huang, Z.J., Zhang, F., Boyden, E.S.*, et al.* (2012). Activation of specific interneurons improves V1 feature selectivity and visual perception. Nature *488*, 379-383.

Li, C.S., Mazzoni, P., and Andersen, R.A. (1999). Effect of reversible inactivation of macaque lateral intraparietal area on visual and memory saccades. J Neurophysiol *81*, 1827-1838.

Li, N., Chen, T.W., Guo, Z.V., Gerfen, C.R., and Svoboda, K. (2015). A motor cortex circuit for motor planning and movement. Nature *519*, 51-56.

Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. Nature *532*, 459-464.

Liu, D., Gu, X., Zhu, J., Zhang, X., Han, Z., Yan, W., Cheng, Q., Hao, J., Fan, H., Hou, R.*, et al.* (2014). Medial prefrontal activity during delay period contributes to learning of a working memory task. Science *346*, 458-463.

McNaughton, B.L., Mizumori, S.J., Barnes, C.A., Leonard, B.J., Marquis, M., and Green, E.J. (1994). Cortical representation of motion during unrestrained spatial navigation in the rat. Cereb Cortex *4*, 27-39.

Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. J Neurosci *16*, 5154-5167.

Mittmann, W., Wallace, D.J., Czubayko, U., Herb, J.T., Schaefer, A.T., Looger, L.L., Denk, W., and Kerr, J.N. (2011). Two-photon calcium imaging of evoked activity from L5 somatosensory neurons in vivo. Nat Neurosci *14*, 1089-1093.

Murakami, M., Vicente, M.I., Costa, G.M., and Mainen, Z.F. (2014). Neural antecedents of self-initiated actions in secondary motor cortex. Nat Neurosci *17*, 1574-1582.

Nakamura, K., and Colby, C.L. (2000). Visual, saccade-related, and cognitive activation of single neurons in monkey extrastriate area V3A. J Neurophysiol *84*, 677-692.

Nitz, D.A. (2006). Tracking route progression in the posterior parietal cortex. Neuron *49*, 747-756.

O'Connor, D.H., Clack, N.G., Huber, D., Komiyama, T., Myers, E.W., and Svoboda, K. (2010). Vibrissa-based object localization in head-fixed mice. J Neurosci *30*, 1947-1967.

Oh, S.W., Harris, J.A., Ng, L., Winslow, B., Cain, N., Mihalas, S., Wang, Q., Lau, C., Kuan, L., Henry, A.M.*, et al.* (2014). A mesoscale connectome of the mouse brain. Nature *508*, 207-214.

Park, I.M., Meister, M.L., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. Nat Neurosci *17*, 1395-1403.

Peron, S.P., Freeman, J., Iyer, V., Guo, C., and Svoboda, K. (2015). A Cellular Resolution Map of Barrel Cortex Activity during Tactile Behavior. Neuron *86*, 783-799.

Pinto, L., and Dan, Y. (2015). Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed Behavior. Neuron *87*, 437-450.

Pinto, L., Goard, M.J., Estandian, D., Xu, M., Kwan, A.C., Lee, S.H., Harrison, T.C., Feng, G., and Dan, Y. (2013). Fast modulation of visual perception by basal forebrain cholinergic neurons. Nat Neurosci *16*, 1857-1863.

Prakash, R., Yizhar, O., Grewe, B., Ramakrishnan, C., Wang, N., Goshen, I., Packer, A.M., Peterka, D.S., Yuste, R., Schnitzer, M.J.*, et al.* (2012). Two-photon optogenetic toolbox for fast inhibition, excitation and bistable modulation. Nature methods *9*, 1171-1179.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat Neurosci *17*, 1784-1792.

Romo, R., Brody, C.D., Hernandez, A., and Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. Nature *399*, 470-473.

Romo, R., and de Lafuente, V. (2013). Conversion of sensory signals into perceptual decisions. Progress in neurobiology *103*, 41-75.

Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. Journal of neurophysiology *86*, 1916-1936.

Sippy, T., Lapray, D., Crochet, S., and Petersen, C.C. (2015). Cell-Type-Specific Sensorimotor Processing in Striatal Projection Neurons during Goal-Directed Behavior. Neuron *88*, 298-305.

Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Coding of intention in the posterior parietal cortex. Nature *386*, 167-170.

Sreenivasan, K.K., Curtis, C.E., and D'Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. Trends Cogn Sci *18*, 82-89.

Super, H., Spekreijse, H., and Lamme, V.A. (2001). A neural correlate of working memory in the monkey primary visual cortex. Science *293*, 120-124.

Wang, Q., Sporns, O., and Burkhalter, A. (2012). Network analysis of corticocortical connections reveals ventral and dorsal processing streams in mouse visual cortex. J Neurosci *32*, 4386-4399.

Wang, X.J. (2008). Decision making in recurrent neuronal circuits. Neuron *60*, 215-234.

Whitlock, J.R., Sutherland, R.J., Witter, M.P., Moser, M.B., and Moser, E.I. (2008). Navigating from hippocampus to parietal cortex. Proceedings of the National Academy of Sciences of the United States of America *105*, 14755-14762.

Zagha, E., Xinxin, G., and McCormick, D.A. (2015). Competing Neural Ensembles in Motor Cortex Gate Goal-Directed Motor Output. Neuron *88*, 565-577.

Znamenskiy, P., and Zador, A.M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. Nature *497*, 482-485.

Zorzos, A.N., Scholvin, J., Boyden, E.S., and Fonstad, C.G. (2012). Three-dimensional multiwaveguide probe array for light delivery to distributed brain circuits. Opt Lett *37*, 4841-4843.

2.8 Figures

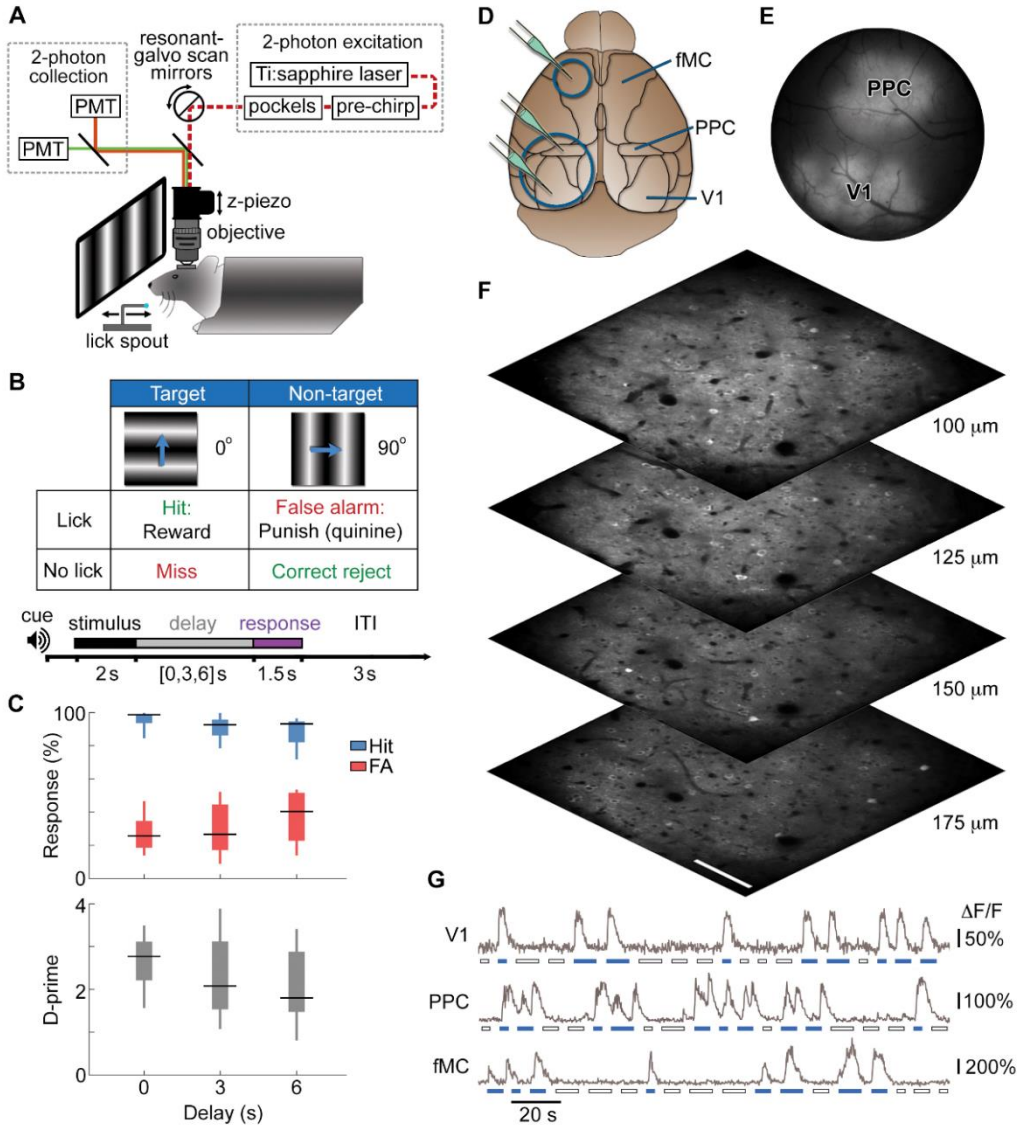

Figure 2.1: Calcium imaging during a sensorimotor decision task.

(**A**) Experimental setup for 2-photon imaging in head-fixed mice performing a memory-guided visual discrimination task. Resonant-galvo scan mirrors were synchronized to a z-piezo to allow volumetric imaging. A retractable lick spout was used to restrict the timing of behavioral responses to a specific epoch of the task. (**B**) Contingency table (top) and trial structure of the memory-guided visual discrimination task (bottom). Trials consisted of stimulus, delay, and response epochs, and the retractable spout was within reach only during the response epoch. Trials of three different delays (0, 3, or 6 seconds) were randomly interleaved. (**C**) Average behavioral performance (n = 8 mice). Top, Response rate for target stimuli (hit rate; blue) and non-target stimuli (false alarm rate; red). Bottom, D-prime for delays of 0-6 seconds. (**D**) Location of cranial windows (blue circles) and AAV-GCaMP6s injections (green pipettes) in primary visual cortex (V1), posterior parietal cortex (PPC), and frontal motor cortex (fMC). (**E**) Example wide-field epifluorescence image of GCaMP6s expression in both V1 and PPC (window diameter, 4 mm). (**F**) Four imaging planes (25 µm apart) within V1 acquired at a stack rate of 5 Hz. Scale bar, 100 µm. (**G**) Sample raw ΔF/F traces from V1, PPC, and fMC (separate experiments) during interleaved target (blue bars) and non-target (white bars) trials of varying delay duration.

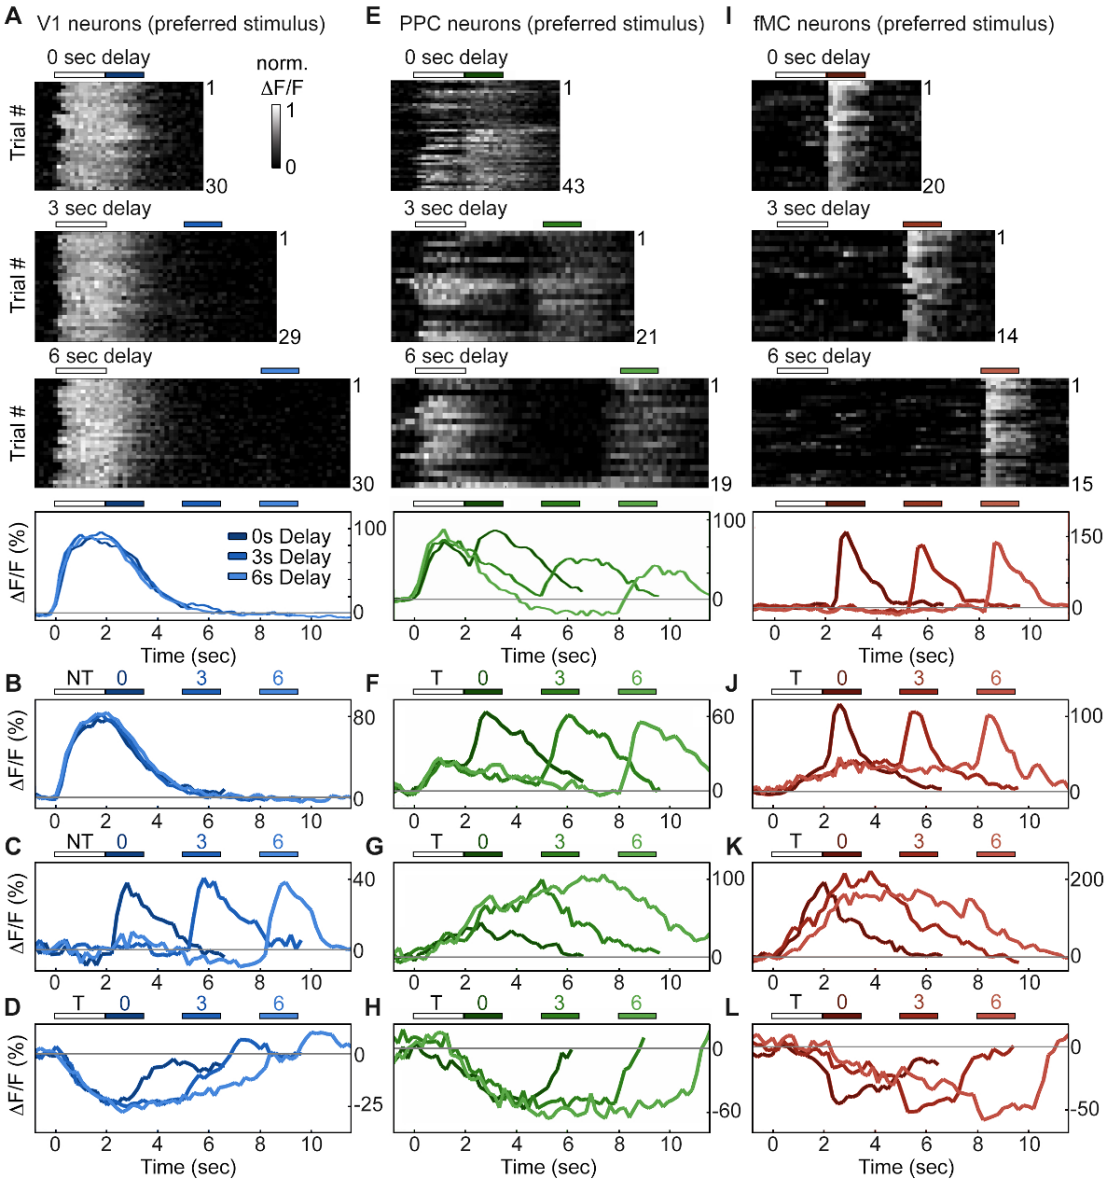

Figure 2.2. Single neurons exhibit heterogeneous trial-evoked responses.

(**A**) Trial-to-trial and average responses from a single V1 neuron. Top three plots, stacked single-trial ΔF/F responses to the preferred stimulus (correct trials only), grouped into 0-, 3-, and 6-second delay trials. The neuron is active during the stimulus for all three delays. Colored bars above plots indicate time of visual stimulus (white) and spout extension (colored, shade indicates delay period). Bottom plot, overlay of mean ΔF/F responses during 0-, 3-, and 6-s delay trials. Shade of average traces indicates delay period (dark blue, 0-s delay; medium blue, 3-s delay, light blue, 6-s delay). Additional V1 example neurons exhibiting activity during the stimulus period (**B**), response period (**C**), or suppressed activity throughout the delay period (**D**). Color shade indicates delay as in (A). (**E**) Same as (A) but for a PPC neuron. This neuron is active during both stimulus and response period. Additional PPC example neurons exhibiting activity during the stimulus and response period (**F**), sustained activity during the delay period (**G**), or suppressed activity throughout the delay period (**H**). Color shade indicates delay as in (E). (**I**) Same as (A) but for a fMC neuron. This neuron is active during the response period. Additional fMC example neurons exhibiting activity during the delay and response period (**J**), during only the delay period (**K**), or suppressed activity throughout the delay period (**L**). Color shade indicates delay as in (I).

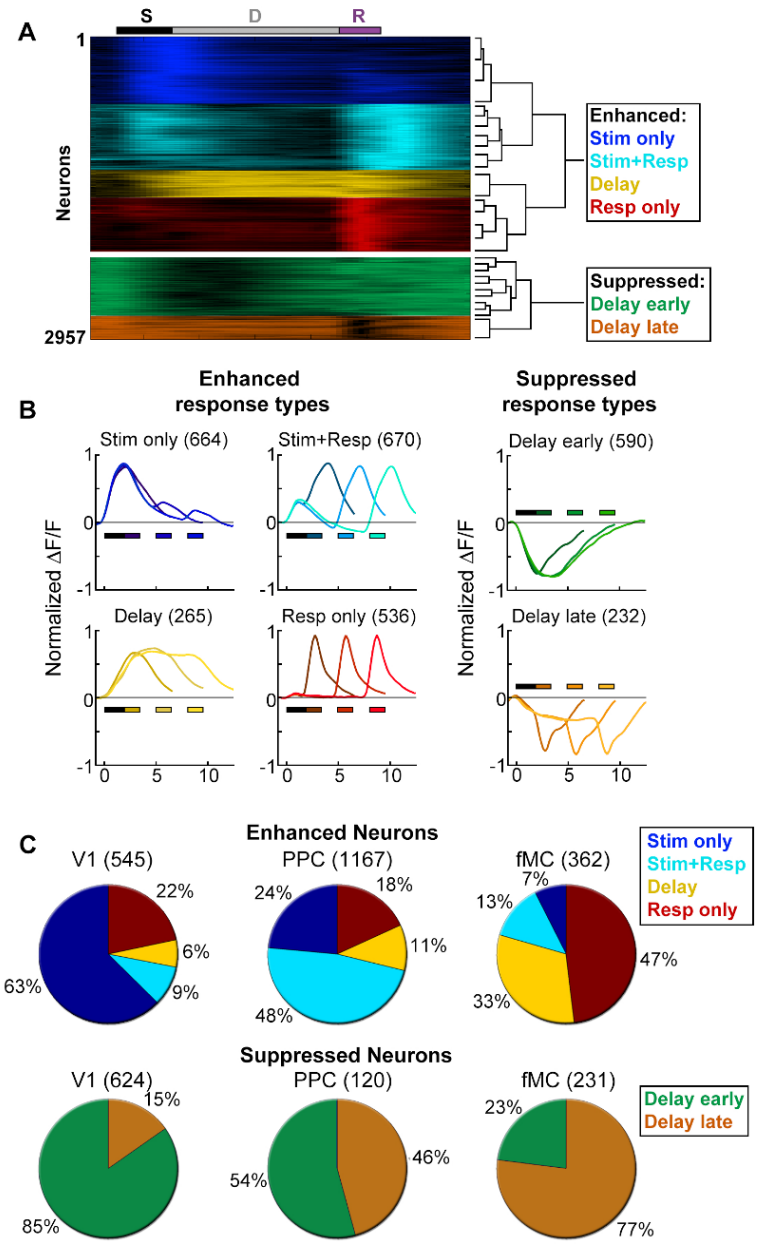

Figure 2.3. Clustering reveals distinct response types distributed between regions.

Hierarchical clustering was performed on the first 20 principal components (explaining >98% of the variance), revealing six separable clusters. Further cluster division resulted in separation of very similar response types. (**A**) Normalized calcium response to the preferred stimulus (correct trials only) of each neuron for the 6 s delay condition, separated by cluster identity (color indicates cluster). Bars at top of plot indicate time of visual stimulus (white), delay (grey), and spout extension (purple). Right inset shows dendrogram resulting from hierarchical clustering procedure and cluster names. (**B**) Average normalized response of each cluster for all delay durations, separated by enhanced and suppressed clusters. Titles indicate name of response type cluster and the number of neurons included. Response color indicates cluster identity. Colored bars indicate time of visual stimulus (white) and spout extension (color indicates cluster, shade indicates delay duration). (**C**) Proportion of neurons in V1 (*n* = 1169), PPC (*n* = 1287) and fMC (*n* = 593) of each response class.

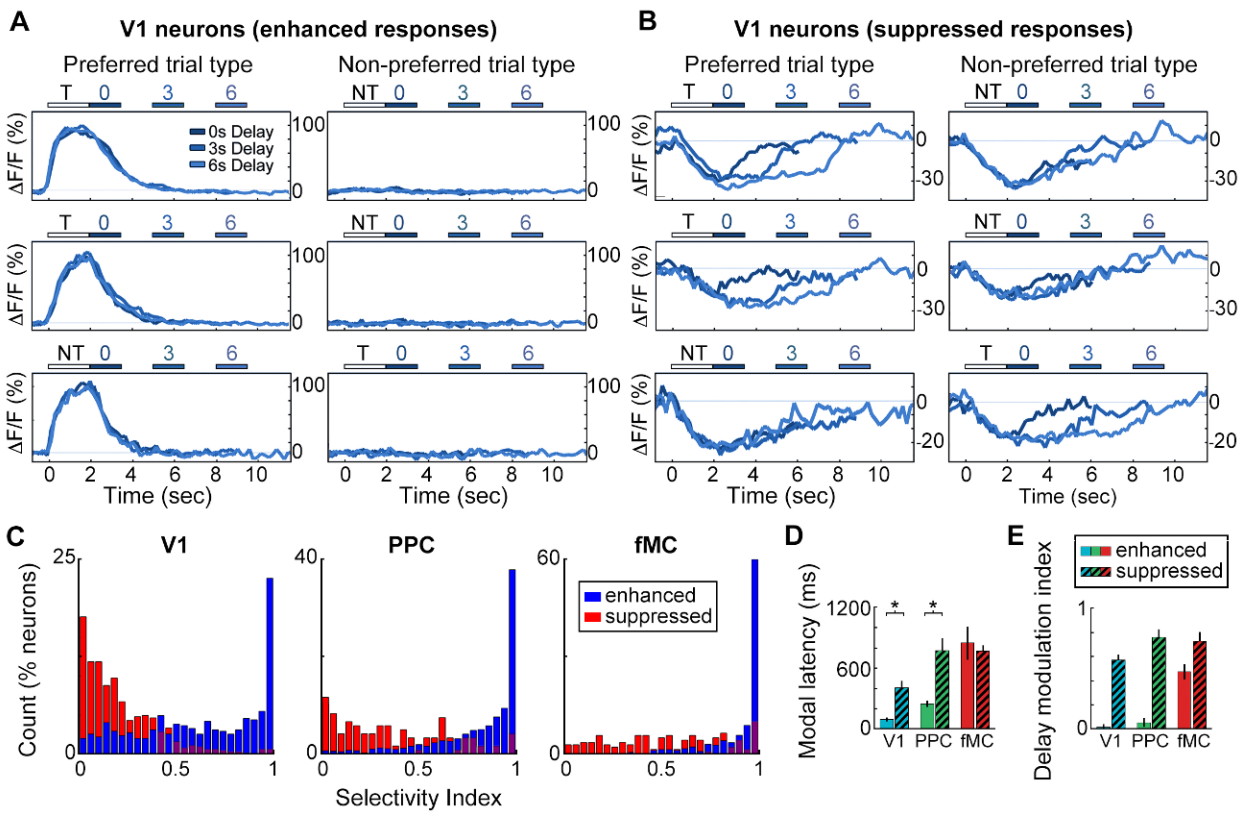

Figure 2.4. Neurons with suppressed activity are much less selective than neurons with enhanced activity.

(**A**) Three example V1 neurons (top, middle, bottom) with enhanced task-driven activity. Responses shown for preferred stimulus (left) and non-preferred stimulus (right), including only correct trials. Each plot shows an overlay of the mean ΔF/F responses during 0-, 3-, and 6-s delay trials. Colored bars above plots indicate time of visual stimulus (white; T, target stimulus; NT, non-target stimulus) and spout extension (blue, shade indicates delay period). Shade of average traces indicates delay period (dark blue, 0-s delay; medium blue, 3-s delay, light blue, 6-s delay). (**B**) Three example V1 neurons (top, middle, bottom) with suppressed task-driven activity. Responses shown for preferred stimulus (left) and non-preferred stimulus (right), including only correct trials. Color shade indicates delay as in (A). (**C**) Histograms of selectivity index for enhanced (blue) and suppressed (red) neurons for each of the three regions (V1, left; PPC, middle; fMC, right). (**D**) Modal latency for enhanced (solid) and suppressed (hashed) neurons for each of the three regions (V1, blue; PPC, green; fMC, red). Latencies for enhanced and suppressed populations differ in V1 and PPC, but not fMC. (**E**) Delay modulation index for enhanced (solid) and suppressed (hashed) neurons for each of the three regions (V1, blue; PPC, green; fMC, red). The delay modulation increases for enhanced neurons from V1 to PPC to fMC, but is high in all three regions for suppressed neurons. Wilcoxon signed-rank test with Bonferroni correction used for all statistical tests. For bar plots, bars indicate mean ± bootstrap-estimated S.E.M.

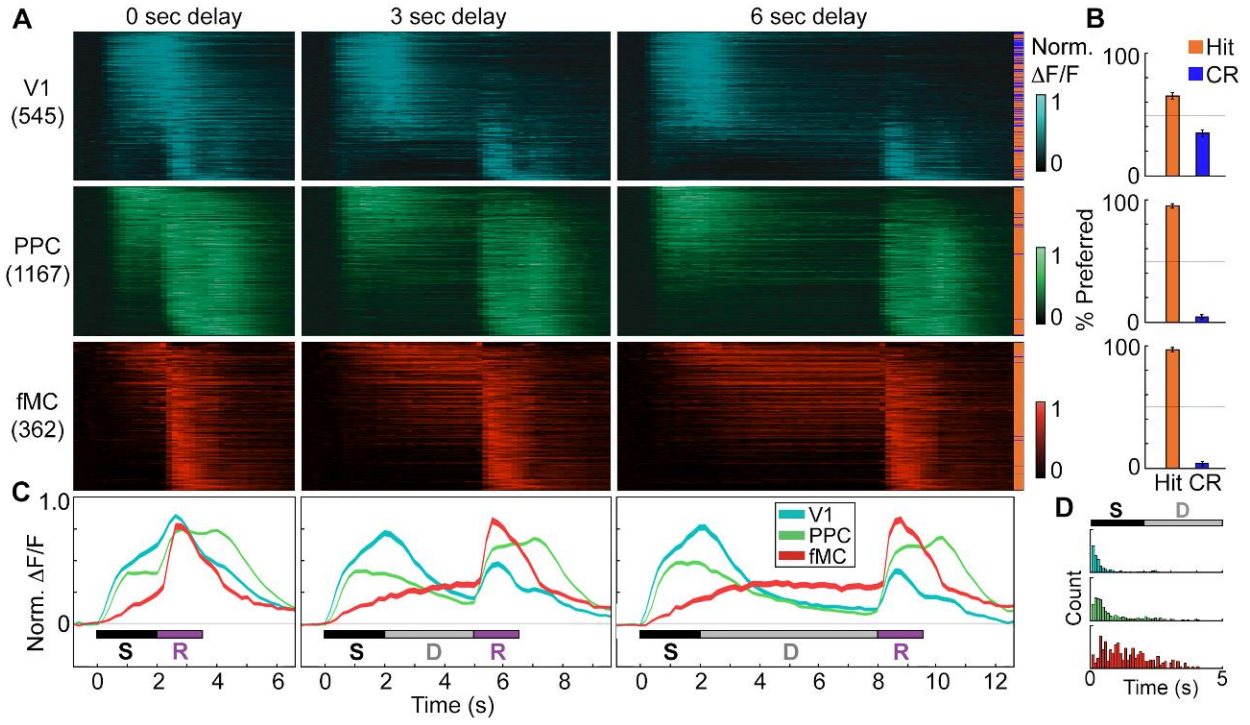

Figure 2.5. Distinct population dynamics in regions V1, PPC, and fMC.

(**A**) Normalized preferred calcium responses of all significantly responsive neurons showing enhanced responses (pooled across all 8 mice) in V1 (top, n = 545), PPC (middle, n = 1167) and fMC (bottom, n = 362), across all three delay durations (0-, 3-, 6-seconds; left, middle, and right, respectively). Only correct trials were included. For each neuron, traces were normalized to the peak of each cell's trial-averaged response (colorbar on right inset). For each area, neurons were sorted by time of peak response. Sidebar indicates preferred trial for each neuron (Hit, orange; CR, blue). (**B**) Proportion of task-responsive neurons in each brain region that preferred Hit (orange) vs. Correct reject (CR, blue) trials. There was a strong bias for neurons in PPC and fMC to toward hit trials. (**C**) Mean population response of each brain region (V1, blue; PPC, green; fMC; red), across the three delay durations (line boundaries indicate mean ± bootstrap-estimated s.e.m.). Colored bars indicate the times of visual stimulus (white), delay (gray), and spout extension (purple). (**D**) Histogram of latency to significant response of each brain region (only neurons with significant responses during stimulus or delay epochs considered). The population of single neuron latencies in V1 significantly precedes PPC, which in turn precedes fMC, for all delays ($p < 10^{-9}$, all comparisons, Wilcoxon rank sum test). Color indicates region (V1, blue; PPC, green; fMC; red).

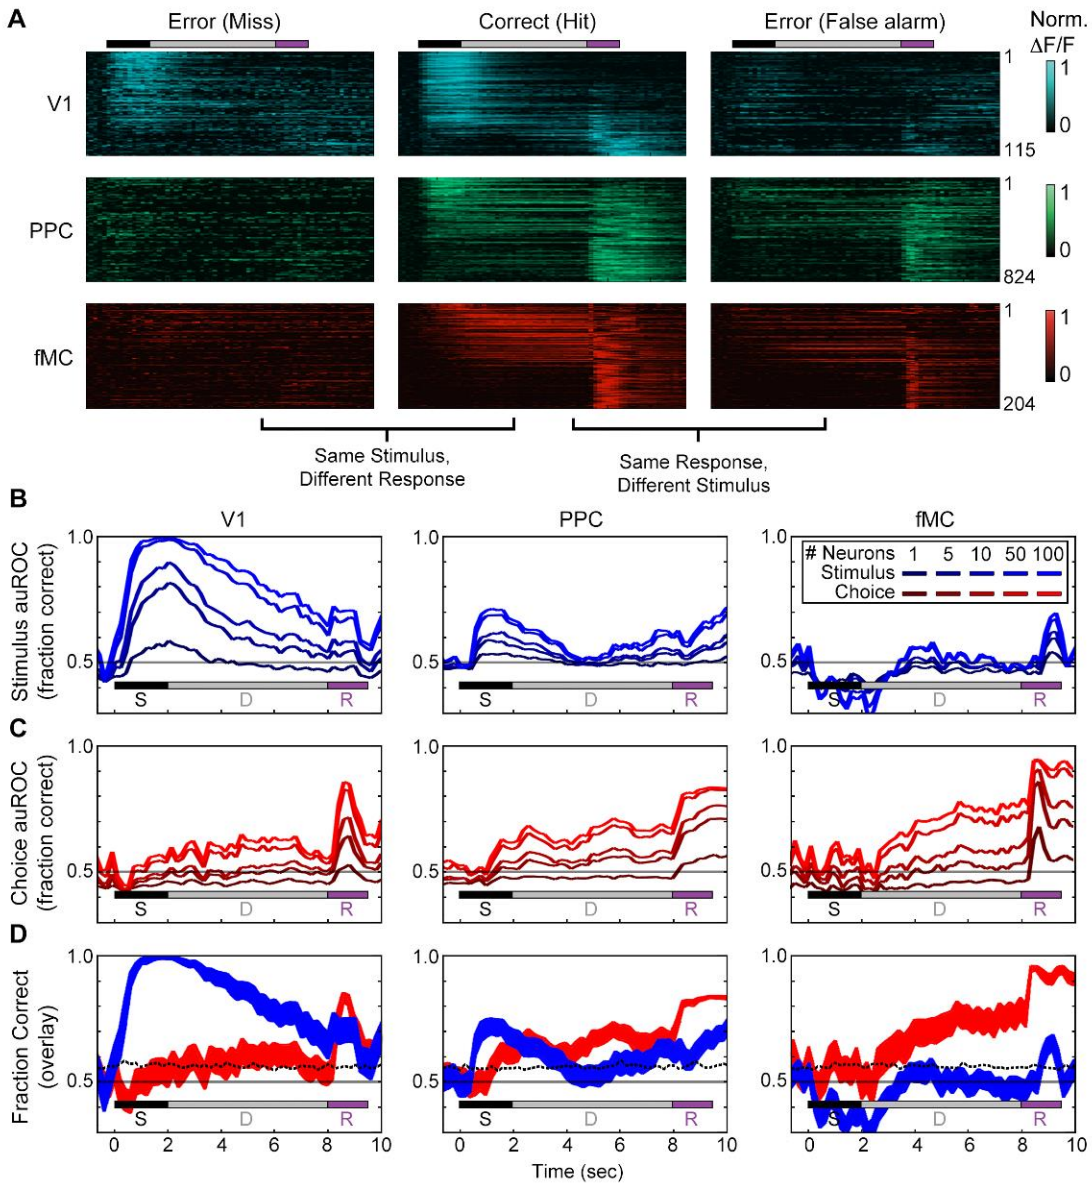

Figure 2.6. Error responses and population coding of task variables.

(**A**) Responses of all target-preferring neurons (6-s delay trials only) in all three regions (V1: top, blue; PPC: middle, green; fMC: bottom, red) to correct target trials (Hit, middle), error trials with the same stimulus but different response (Miss, left), and error trials with the same response but different stimulus (False alarm, right). Color bars indicate normalized $\Delta F/F$ (0 to 1). (**B**) Ideal observer performance (area under the receiver operator characteristic, auROC) in estimating the stimulus identity (target, non-target) based only on the neural activity of significantly enhanced neurons in V1 (left), PPC (middle), and fMC (right), including randomly sampled correct and error trials). Line shading indicates neuron population size used for the decoding (dark to light blue: n = 1 to 100 neurons; see legend). (**C**) Ideal observer performance in estimating the behavioral choice (lick, no-lick) based only on the neural activity of significantly enhanced neurons in V1 (left), PPC (middle), and fMC (right), including randomly sampled correct and error trials). Line shading indicates neuron population size used for the decoding (dark to light red: n = 1 to 100 neurons). (**D**) Re-plot of stimulus identity and behavioral choice decoding performance for n = 100 neurons, with line boundaries indicating mean ± bootstrap-estimated S.E.M. Dotted lines indicate upper 95% confidence interval for 1000 permutations with shuffled labels.

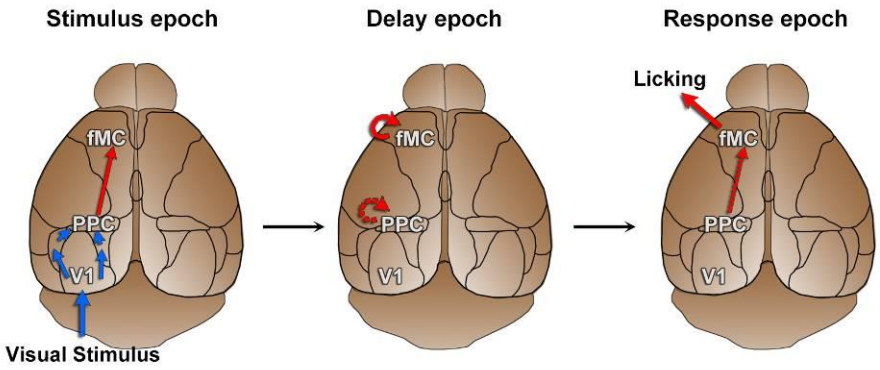

Figure 2.7. Working hypothesis of information propagation.

Based on imaging results (**Figure 2.5**) and population analyses (**Figure 2.6**), we developed a working hypothesis of information propagation during the sensorimotor decision task. During the stimulus epoch (left), stimulus identity (blue) is encoded by V1 and PPC, while behavioral choice (red) is encoded by PPC and fMC. Note that although there are direct inputs from V1 to PPC and PPC to fMC, there may be intermediate connections as well. During the delay epoch (middle), coding of stimulus identity fades while coding of behavioral choice is actively maintained in fMC (and possibly PPC). During the response epoch, the stored choice information leads to the execution of the motor action within fMC (and possibly V1 and PPC).

Chapter 3: Causal roles of visual, parietal, and frontal motor cortex in a sensorimotor decision[1].

3.1 Summary

Neural correlates of sensorimotor decisions have been measured across many brain areas, but the causal role of these regions is poorly understood. We investigated the necessity of visual (V1), posterior parietal (PPC), and frontal motor (fMC) cortices during distinct epochs of a memory-guided visual discrimination task using optogenetic inactivation. Bilateral, reversible silencing of any of the three areas during the stimulus epoch disrupted performance, even using a brief (250 ms) stimulus presentation. However, only fMC was required during the delay and choice epochs. These results lead to the hypothesis that stimulus identity is rapidly transformed into behavioral choice, requiring V1, PPC, and fMC during the transformation period, but only fMC for maintaining the choice in memory prior to execution. However, simultaneous recording and inactivation experiments additionally reveal that disruption of PPC indirectly affects V1 activity, raising the possibility that permissive feedback inputs are critical for behavioral performance.

3.2 Introduction

Perceptual decisions involve multiple cognitive capacities including sensory processing, evidence accumulation, sensorimotor transformation, short-term memory maintenance, movement planning, and motor execution. It is no wonder that neural correlates of short-term memory and perceptual decision-making have been recorded in a plethora of brain regions, including prefrontal (Fuster and Alexander, 1971; Kim and Shadlen, 1999; Kojima and Goldman-Rakic, 1982; Romo et al., 1999), parietal (Constantinidis and Steinmetz, 1996; Shadlen and Newsome, 2001), sensory (Super et al., 2001), and motor (de Lafuente and Romo, 2005; di Pellegrino and Wise,

[1]Many of the findings (Figures 3.1-3.4) presented in this chapter appeared in Goard, Pho et al., 2016

1993) cortices, as well as in several subcortical areas including the striatum (Ding and Gold, 2010) and superior colliculus (Horwitz and Newsome, 1999). Indeed, a few impressive studies and research programs have measured from many of these regions in the same decision task (Brody and Hanks, 2016; Goard et al., 2016; Hernandez et al., 2010; Siegel et al., 2015), and found different neural signatures that appear to vary with each area's position on the sensory to motor hierarchy. However, despite a wealth of research in this area, it remains unclear which areas are causally necessary for the transformation and maintenance of task information.

For example, single-unit electrophysiological recordings in non-human primates have long implicated the posterior parietal cortex (PPC) as a critical node for the process of sensorimotor transformation (Andersen and Cui, 2009; Bennur and Gold, 2011; Freedman and Assad, 2006; Gold and Shadlen, 2007; Shadlen and Newsome, 2001). However, recent inactivation studies in macaques (Katz et al., 2016) and in rodents (Erlich et al., 2015; Guo et al., 2014) have challenged this perspective, finding minimal role for PPC in certain decision tasks. Other studies suggest that motor regions, whether cortical (Guo et al., 2014; Murakami et al., 2014; Zagha et al., 2015) or subcortical (Kopec et al., 2015; Znamenskiy and Zador, 2013), may instead be the key site of sensorimotor transformation. The discrepancies between these different studies may include differences in species, sensory modality, or behavioral task.

Similarly, the importance of various areas in maintaining information in memory remains unresolved. Sustained activity during delayed response tasks have been observed in many areas including prefrontal cortex (Funahashi et al., 1989; Fuster and Alexander, 1971; Kojima and Goldman-Rakic, 1982; Miller et al., 1996; Romo et al., 1999; Sreenivasan et al., 2014), parietal cortex (Chafee and Goldman-Rakic, 1998; Constantinidis and Steinmetz, 1996; Harvey et al., 2012; Shadlen and Newsome, 2001; Snyder et al., 1997), motor cortex (de Lafuente and Romo, 2005; di Pellegrino and Wise, 1993; Erlich et al., 2015; Guo et al., 2014; Hernandez et al., 2010; Li et al., 2015), and even primary sensory areas (Nakamura and Colby, 2000; Super et al., 2001).

However, recent studies in rodents have indicated that some of the regions traditionally thought to be crucial for short-term memory maintenance, such as parietal and prefrontal cortices, appear not to play a major role in some tasks (Erlich et al., 2015; Guo et al., 2014; Liu et al., 2014), though other studies have found an important role for sustained activity in motor regions consistent with theoretical models (Kopec et al., 2015; Li et al., 2016).

To further investigate these questions of sensorimotor transformation and short-term memory maintenance, we hoped in this work to take a more comprehensive approach toward measuring and perturbing activity across sensory, parietal association, and motor cortical regions during a delayed-response task. To leverage the genetic tractability of the mouse toward a better mechanistic understanding of memory-guided sensorimotor decisions, we modified a visual discrimination task used for head-fixed mice (Andermann et al., 2010) by using a retractable spout to separate the components of the task into discrete stimulus, delay, and response epochs. We took advantage of high sensitivity genetically-encoded calcium indicators (Chen et al., 2013b) to enable high-yield 2-photon volume scanning approach (Kampa et al., 2011). This allowed us to image from large populations of neurons in sensory, association and frontal motor regions during discrete epochs of a memory-guided sensorimotor decision task, including stimulus encoding, delay and behavioral response (see **Chapter 2**).

Although measurement of neural activity is an important first step toward defining task-related regions, the presence of neural activity does not prove that a given region plays a causal role in mediating behavior. For example, it is possible that sustained activity seen in one region reflects sustained activity in a connected region without playing a direct role in the maintenance of mnemonic information. Microstimulation studies have revealed that activation of neurons in certain regions is sufficient to influence task performance (Bisley et al., 2001; Hanks et al., 2006), but does not directly address whether the region is necessary. Other studies have used surgical lesions (Gisquet-Verrier and Delatour, 2006; Sakurai and Sugimoto, 1985), pharmacological

inactivation (Erlich et al., 2015; Harvey et al., 2012; Katz et al., 2016; Li et al., 1999) or tissue cooling (Bauer and Fuster, 1976; Chafee and Goldman-Rakic, 2000) to reveal the cortical regions necessary for performance of short-term memory maintenance in rodents and non-human primates. However, while these studies have been important for establishing the anatomical framework of memory maintenance, these manipulations have limited spatial and temporal resolution. In particular, activity could not be silenced rapidly enough to test specific epochs of the task (e.g., delay period), and it is thus unclear whether a given manipulation affects stimulus perception, memory maintenance, or motor output.

Recent optical inactivation approaches have revealed that the effect of cortical inactivation on behavior is crucially dependent on timing (Kopec et al., 2015; Li et al., 2016; Sachidhanandam et al., 2013) and whether inactivation is unilateral or bilateral (Li et al., 2016). To continue this approach, we used an optogenetic approach for reversibly silencing activity bilaterally in defined cortical regions with precise temporal control (Zhao et al., 2011). We modified a visual discrimination task used for head-fixed mice (Andermann et al., 2010) by using a retractable spout to separate the components of the task into discrete stimulus, delay, and response epochs. Using inactivation of bilateral cortical regions exhibiting task-related responses, we sought to determine the necessity of sensory, association, and frontal motor cortical regions during each epoch of a memory-guided task.

One major caveat, however, of any perturbation approach is the possibility of off-target effects. Although the reductionist approach of isolating specific roles to individual brain areas is attractive for its conceptual simplicity, the reality is that the brain is one densely interconnected dynamical system. Sudden and acute perturbations of one node may have non-physiological and compromising effects on other nodes in the circuit, whether by disrupting downstream dynamics or removing permissive input that disrupts a circuit's excitatory-inhibitory balance. Indeed, as recent work has demonstrated in both rodents and songbirds, acute manipulation of brain regions

can disrupt behavior even when those same brain regions are dispensable after chronic lesion (Otchy et al., 2015). This effect arises because transient manipulations affect downstream circuits which are essential for behavior (as shown by chronic lesion). Transient manipulations alone therefore cannot distinguish whether an area truly plays an instructive role in behavior and provides necessary computations for task performance (Otchy et al., 2015; Panzeri et al., 2017), or whether an area plays merely a permissive role in providing necessary input to other essential regions.

Distinguishing these alternatives is particularly important for the study of the posterior parietal cortex in rodents. Recent inactivation studies observed a minimal role for PPC in rat auditory-based (Erlich et al., 2015) and mouse whisker-based (Guo et al., 2014) decision tasks. However, PPC does appear to be necessary for visual sensorimotor decision tasks in mice (Harvey et al., 2012) and in rats (Licata et al., 2016; Raposo et al., 2014). Some have argued that the proximity of PPC to secondary visual regions makes it difficult to clearly isolate the effect of PPC inactivation (Brody and Hanks, 2016; Erlich et al., 2015). Indeed, both anatomical projection studies (Wang and Burkhalter, 2007; Wang et al., 2012), as well as functional mapping studies (Garrett et al., 2014; Marshel et al., 2011) indicate that PPC lies immediately adjacent to, or perhaps overlaps with, a group of retinotopically-organized extrastriate areas that are rostral to V1. The fact that these visual areas are very small, and are all directly and indirectly interconnected to one another (Wang et al., 2012) means that great care must be taken to properly interpret their individual roles in behavior.

To isolate the role of PPC in visual sensorimotor decisions, we performed control experiments to simultaneously record from visual cortex while manipulating activity in PPC. By isolating the effects of light stimulation on different putative cell types, we were able to distinguish direct spread of light from indirect network-level effects. We find that PPC manipulation

moderately suppresses visual cortex activity, likely via network mechanisms, raising the possibility that feedback from PPC to V1 plays a permissive role in behavior.

3.3 Experimental Procedures

3.3.1 Surgical procedures

All procedures were approved by the Massachusetts Institute of Technology Animal Care and Use Committee. Data were collected from male adult (60-120 day old) VGAT-ChR2-EYFP mice (Jackson Laboratory; n = 15). The animals were housed on a 12 hour light/dark cycle in cages of up to 5 animals before the implants, and individually after the implants. All surgeries were conducted under isoflurane anesthesia (3.5% induction, 1.5-2.5% maintenance). Meloxicam (1 mg kg^{-1}, subcutaneous) was administered pre-operatively and every 24 hours for 3 days to reduce inflammation. Surgical procedures for the first headplate implant surgery are identical to those described in Section 2.3.1. After head plate implantation, mice recovered for at least five days before beginning water restriction.

After behavioral training was complete, animals were taken off water restriction for five days before undergoing a second surgery to implant the photoinhibition windows. Procedures for anesthetic administration and post-operative care were identical to the first surgery. The dental acrylic and silicon elastomer covering the targeted region were removed using a drill burr. The skull surface was then cleaned and a craniotomy (2-4 mm, depending on targeted structure) was made over the region of interest, leaving the dura intact. Craniotomies were made bilaterally, with no virus injection. Finally, a cranial window was implanted over the craniotomy and sealed first with silicon elastomer then with dental acrylic. The cranial windows were made of two rounded pieces of coverglass (Warner Instruments) bonded with optical glue (NOA 61, Norland). The bottom piece was circular or oval, custom cut according to cortical region(s) (V1: 2.5 mm x 2.5

mm; PPC: 1 mm x 2 mm; V1 + PPC: 4 mm x 4 mm; fMC: 2 mm x 2.5 mm, anterior-posterior x medial-lateral) and fit snugly in the craniotomy. The top piece was a larger circular coverglass (3-5 mm, depending on size of bottom piece) and was bonded to the skull using dental acrylic. Mice recovered for five days before commencing water restriction.

3.3.2 Behavior

The behavioral apparatus and training procedures are identical to those described in Section 2.3.2. Briefly, mice were head-fixed using optical hardware (Thorlabs) and placed in a polypropylene tube to limit movement. Spout position was controlled by mounting the spout apparatus on a pressure-driven sliding linear actuator (Festo) controlled by two solenoids (Parker). Licks were detected using an infrared emitter/receiver pair (Digikey) mounted on either side of the retractable lick spout. Rewards consisted of 5-8 μl water and punishments consisted of a white noise auditory stimulus alone (early training) or white noise plus 1-3 μl of 5mM quinine hydrochloride (Sigma) in water (late training). The stimulus consisted of sine wave gratings (spatial frequency: 0.05 cycles deg^{-1}; temporal frequency: 2 Hz) drifting at either 0 degrees (target) or 90 degrees (non-target) away from vertical.

Mice were trained in successive stages, with advancement to the next stage contingent on correct performance. At first the spout was extended within range before the stimulus appeared, then spout extend time was gradually delayed until after the stimulus had turned off (i.e., 0 s delay). Finally, the variable delay period was gradually increased to 0/2/4 s.

Once mice reached high levels of performance at the final stage of the task ($d' > 1.5$ and $R_{HIT} - R_{FA} > 50\%$), they were removed from water restriction for window implantation. After recovery from window implantation surgery, they were re-trained to a level of high performance

(2-7 days) before beginning experimental sessions. Any sessions with poor performance were discarded (minimum performance criterion: $d' > 1$ and $R_{HIT} - R_{FA} > 30\%$).

3.3.3 Cell-attached recordings

For *in vivo* anesthetized electrophysiology experiments (**Figure 3.1A,B**), we used procedures similar to those previously described (Wilson et al., 2012). Briefly, transgenic VGAT-ChR2-EYFP mice (n = 2) were anesthetized with isoflurane (1.5%), with body temperature maintained at 37.5º C with a heating blanket. A metal head plate was attached to the skull using superglue and dental acrylic, and a 1mm craniotomy was performed over V1. The dosage of anesthesia was then lowered (1%) before beginning electrophysiology.

Recordings were made using custom software written in Matlab (Mathworks) controlling a MultiClamp 700B Amplifier (Axon) that measured differences between a glass pipette electrode inserted into the brain at 27º and an Ag/AgCl ground pellet electrode positioned in the same solution as the brain. Borosilicate pipettes (outer diameter = 1.5mm, inner diameter = 1.17mm) were pulled using a Sutter P-2000 laser puller (Sutter Instruments) to a diameter corresponding to 3-7 MΩ. The pipette was back-filled with Alexa Fluor 488 (Molecular Probes) and targeted to the injection site using a 10x lens. Cells were targeted blindly by advancing diagonally through the cortex with light positive pressure. Cell proximity was detected through deflections in electrical resistance observed in voltage clamp during a repetitive 5mV command pulse. Once resistance had increased a few MΩ, slight negative pressure was applied and the pipette was advanced more slowly until resistance further increased (to a final value of 10-30 MΩ) and/or spikes were detected visually or via an audio monitor. At that point, the amplifier was switched to current clamp to record spikes.

Optogenetic activation of local inhibitory cells was achieved using a 200mW 473nm diode-pumped solid state blue laser (Opto-Engine) coupled with a 200μm fiber. Laser intensity was modulated with a variable neutral density filter (Thorlabs) to match the intensity used during behavioral experiments (6.5 mW mm^{-2}). Full-field drifting grating stimuli were presented for 4 seconds, with each presentation preceded by a 6 second “off” period with gray screen. Optogenetic stimulation occurred on every other stimulus presentation during the middle 2 seconds of the “on” period. Data was analyzed using Matlab, with spikes detected using a manually-defined thresholds.

3.3.4 Photoinhibition

Blue light illumination was provided by a 200 mW 473 nm diode-pumped solid state laser (MBL-III-473, Opto-Engine). The laser was passed through a 50/50 splitter (CM1-BS013, Thorlabs), with each output passed through an adjustable neutral density filter (Thorlabs) into a fiber launch (PAF-X-11-PC-A, Thorlabs). Fibers terminated in ceramic ferrules that were precisely positioned above the cranial window with optical hardware. A light shield (eMachineShop.com) was attached to the head plate to prevent reflected laser light from reaching the retina and influencing behavior. Laser power was controlled by analog inputs sent from the behavior computer. Continuous light pulses (2.2 s, 6.5 mW mm^{-2}) were used to suppress activity in a sustained manner at low laser power (Zhao et al., 2011). We measured the spread of inhibition using immunohistological labeling of activity-driven protein expression (see Histology section). For photoinhibition experiments, laser stimulation was given on 50% of trials with 4 s delay (25% of all trials; no photoinhibition on 0 s or 2 s delay trials) in a randomly inter leaved manner. Laser stimulation was applied throughout the duration of visual stimulus presentation (stimulus epoch), in the middle 2 s of delay (delay epoch), or during the spout extension (response epoch).

3.3.5 Histology

For photoinhibition characterization experiments (**Figure 3.1C-E**), naïve VGAT-ChR2-EYFP or wild-type mice were implanted with a head plate and a 2 mm diameter cranial window over posterior cortex (V1 or PPC). After recovery and habituation, mice were head-fixed and stimulated with 2 s light pulses (6.5 mW mm^{-2}) every 10 s. Drifting grating stimuli were played during the light pulses to drive neural activity. After 1.5 h of light stimulation, mice were removed from the rig and perfused with 4% paraformaldehyde. The brain was removed and placed in 4% PFA for 24 h, followed by 24 h in PBS. 50 μm sagittal slices were cut and placed in PBST (0.2% Triton X-100 in PBS) with 5% normal goat serum for 1 h. They were then incubated with primary antibody at 4°C for 24 h (rabbit Anti-c-Fos 1:200, SC-52, Santa Cruz Biotechnology, Inc.; rabbit Anti-Arc 1:200, ab118929, Abcam). After washing in PBST, slices were incubated 1 h with secondary antibody (1:200 AlexFlour555 anti-rabbit, Invitrogen). Slices were washed again in PBST, mounted and coverslipped. Slices were imaged on a laser scanning confocal microscope (Zeiss). Repeating this procedure with Arc as an activity marker (1:500 anti-Arc, Abcam) yielded qualitatively similar results as c-Fos (data not shown).

3.3.6 Silicon probe recordings

In vivo single unit recordings (**Figure 3.6**, **Figure 3.7**, **Figure 3.8**) were performed in untrained *VGAT-ChR2-EYFP* mice (n=6) under isoflurane anesthesia (1.5%) with body temperature maintained at 37.5º C with a heating blanket. A metal headplate was attached to the skull using dental acrylic, and two craniotomies were performed on the left hemisphere: one over PPC (1 mm x 2mm, AP x ML) and one over V1 (1.5 mm x 1.5mm), centered at the same coordinates described above. Dental acrylic mixed with black ink was applied to the area of skull between the two

craniotomies, to reduce light transmission. The dosage of anesthesia was then lowered (0.8-1%) before beginning electrophysiology.

Recordings were made using a silicon linear microprobe with 16 sites spaced at 50µm intervals (NeuroNexus, model A1x16-5mm-50-177-A16). The electrode was inserted at a 45º degree angle, and was moved using a micromanipulator (Sutter, MP-285) to a depth of at least 700µm to ensure that all sites were within the cortex. The electrode was allowed to settle for at least 5 minutes to ensure stability of recordings. Electrical signals, referenced to a chlorinated silver ground wire that was placed in the same saline solution as the brain, were amplified and recorded using a Plexon system (16 channels, 50kHz at 12-bit resolution). The presence of spiking activity was assessed online using an audio monitor, and channels with spiking activity were selected for post-hoc offline spike sorting (Offline Sorter 2.8.8, Plexon).

Prior to the first recording of each experiment, two optical fibers terminating in ceramic ferrules were precisely positioned above the two craniotomies (V1 and PPC) to form a 2.5mm x 2.5mm spot over each. As in the behavioral experiments, this spot exceeded the size of the PPC craniotomy, but provided a constant light density of 6.5 mW mm^{-2}. Only one fiber was connected to the laser (473 nm, MBL-III-473, Opto Engine) at a time, and the connection was manually switched in blocks of trials for each recording.

Visual stimuli were presented to the right eye alone using a screen placed at an oblique angle to the animal. The stimulus consisted of sinusoidal drifting gratings (0.05 cycles deg^{-1}; 2 Hz) of 4 different orientations, all presented within 2 s (0.5 s per orientation). This was chosen to activate a broad population of V1 neurons irrespective of tuning. Stimulus presentation was interleaved with an 8 s gray screen period. On every other stimulus presentation, laser stimulation was also applied for 2.2 s, beginning 0.1 s before the onset of the visual stimulus. Laser stimulation was applied to PPC or V1 in blocks of 30-50 trials. In most recordings, a block of trials with low-power (total power: 19 mW, density: 3.9 mW mm^{-2}) PPC stimulation was also recorded.

Offline analysis to sort waveforms for each unit was performed using commercial programs (Offline Sorter 2.8.8, Plexon). Briefly, raw voltage traces were high-pass filtered at 250Hz, events were detected by voltage threshold crossing, and clusters were separated in principal components space. Post-processing was performed using custom MATLAB code. Units were included for further analysis if they exhibited clear isolation from noise and other units, robust and stable firing across at least two blocks of trials, and visual responsiveness (t-test, $p < 0.05$). Neurons were further separated into two groups based on the effect of V1 stimulation on their responses.

Peri-stimulus time histograms were generated using 0.1 ms time bins. Spike widths were calculated by upsampling the average waveform to 500kHz (2μs resolution) by linear interpolation, and measuring the time between the peak and trough. Normalized spike rate was calculated within each block as the firing rate on laser trials divided by the firing rate on no-laser trials. Normalized spike rates below 0.1 were rounded down to 0 for logarithmic plots. Distances from the edge of laser spot over PPC were estimated using micromanipulator measurements and known site locations along the electrode. The position of the electrode tip was noted for each recording, and the location of each unit along the probe was inferred by assuming a 45-degree angle and calculating the X-Y Euclidean distance (ignoring depth) from the edge of the PPC window. Units were recorded from a range of depths and cortical layers (range: 160-784 μm, median: 458 μm).

3.3.7 General statistics

Data groups were tested for normality using the Kolmogorov-Smirnov test and then compared with the appropriate tests (t-tests, Wilcoxon rank-sum or Wilcoxon signed-rank tests, all two sided). Bonferroni correction was used for multiple comparisons. To ensure that results were

replicable between sessions and mice, we included 4-6 mice for each region for the inactivation experiments.

3.4 Results

3.4.1 Summary of behavioral task, imaging results, and working hypothesis

As previously described in Chapter 2, we trained head-fixed mice to perform a visual discrimination task with a memory-guided response (see **Figure 2.1**). In this task, water-restricted mice were presented a 2 s drifting grating stimulus at one of two orientations (0°, 90° from vertical), followed by a delay period, at which point a lick spout was moved rapidly into reach with a linear actuator for 1.5 s. Licking on "go" trials (horizontal grating drifting toward 0° from vertical) was rewarded with 5-8 μl water (hit), while licking on "no-go" trials (vertical grating drifting toward 90° from vertical) was punished with 2 μl water containing 5mM quinine hydrochloride (false alarm). This structure allowed the separation of each trial into "stimulus", "delay", and "response" epochs.

We focused our experiments on three cortical regions we expected to be important for performance of this task: (1) the primary visual cortex (V1), which is known to be important for orientation discrimination (Glickfeld et al., 2013); (2) the posterior parietal cortex (PPC), which receives extensive input from visual regions (Harvey et al., 2012; Oh et al., 2014; Pho et al., 2015), projects to motor regions (Wang et al., 2012), and has been implicated in sensorimotor decision tasks (Andersen and Cui, 2009; Gold and Shadlen, 2007; Hanks et al., 2015; McNaughton et al., 1994; Nitz, 2006; Raposo et al., 2014; Shadlen and Newsome, 2001; Whitlock et al., 2008) and (3) the frontal motor cortices (fMC), which include regions known to be crucial for voluntary licking behaviors (Guo et al., 2014; Komiyama et al., 2010).

As shown in detail in Chapter 2, population calcium imaging during behavior revealed that all three regions exhibit neural responses during all three task epochs, with a surprising amount

of heterogeneity in each region across individual neurons (see **Figures 2.2-2.4**). However, population coding analyses revealed that the three brain regions play distinct roles in encoding the task. V1 principally encodes the stimulus identity, PPC encodes a mixture of stimulus identity and behavioral choice, and fMC predominantly encodes the behavioral choice throughout the stimulus, delay, and response epochs (see **Figures 2.5-2.6**). Together, these results suggest a working hypothesis of the roles that neural activity in regions V1, PPC, and fMC play in the performance of the task. Specifically, the results suggest that sensory input is primarily processed during the stimulus epoch, first in V1, then subsequently in PPC. Neural activity related to the behavioral choice arises in PPC and fMC shortly after the peak in stimulus identity coding, and is sustained in both regions throughout the delay and response epochs. This implies that stimulus identity is rapidly transformed into a behavioral choice within the stimulus epoch (possibly within PPC), and then the behavioral choice is maintained in higher regions (potentially in both PPC and fMC) until the relevant motor action is performed (see **Figure 2.7**).

3.4.2 Characterization of photoinhibition approach

Although trial-locked neural activity provides insight into which cortical regions may contribute to behavioral performance during different epochs of the task, it is important to note that the presence of neural activity in a region does not demonstrate a functional role in task performance. To directly test the plausibility of our working hypothesis (see Figure 2.7), we next used an optogenetic inactivation approach to directly test the necessity of spatially defined cortical regions during specific epochs of the task. To inactivate the cortical regions of interest, we utilized the VGAT-ChR2-EYFP transgenic mouse line (Zhao et al., 2011) to rapidly and reversibly inhibit neural activity. The VGAT-ChR2-EYFP mice express the light-sensitive cation channel channelrhodopsin (ChR2) in all cortical inhibitory neuron subtypes, and thus stimulation of cortex with 473nm light effectively silences the activity of nearby pyramidal cells (Guo et al., 2014; Zhao

et al., 2011). To characterize the efficacy of optogenetic inhibition, we performed cell-attached recordings from regular-spiking neurons in L2/3 and L5 of region V1 (**Figure 3.1A**). We found that laser stimulation (2 s continuous pulses, 6.5mW mm^{-2}) potently inhibited the activity of regular spiking neurons (see example L5 neuron, **Figure 3.1B**), completely eliminating spikes during the laser stimulation period in all recorded neurons, even during presentation of a drifting grating stimulus of the preferred orientation (grating presented from 1 s before laser stimulation to 1 s after laser stimulation; **Figure 3.1B**, inset, n = 2 mice). To characterize the spatial extent of photoinhibition, we performed immunohistochemistry using c-Fos as a marker of neural activity (**Figure 3.1C**). Naïve VGAT-ChR2-EYFP and control mice were implanted with windows and periodically illuminated with 2s continuous pulses of 473nm light (6.5mW mm^{-2}). We found that after light exposure, VGAT-ChR2-EYFP mice (but not wild-type controls) showed dramatically reduced expression of c-Fos throughout all layers of the cortex underlying our cranial window (**Figure 3.1D,E,** n = 3 mice). The suppression was confined to within 500 μm laterally from the edge of the window, and did not spread past the white matter into subcortical regions. Similar results were achieved using an Arc antibody (n = 2 mice).

Using a simplified version of the task with a 0 s delay period, we confirmed that bilateral inactivation of V1 during the stimulus epoch impaired performance of VGAT-ChR2-EYFP mice ($\Delta d'$ = *-1.41,* p = 0.002), but not wild-type controls exposed to the same light stimulation ($\Delta d'$ = +0.14, p = 0.27; **Figure 3.2A-C**). In some neurons, we noticed that it took a few hundred milliseconds for the spike rate to recover to baseline (e.g., neuron in **Figure 3.1B**). To ensure that inactivation did not have long-lasting effects on the inactivated cortical regions, we inactivated region V1 for 2 s prior to the stimulus onset, and found no effect on performance ($\Delta d'$ = -0.15, p = 0.51); **Figure 3.2D-F**).

3.4.3 Photoinhibition of specific cortical regions during task performance

After training VGAT-ChR2-EYFP mice on the task, we implanted bilateral glass cranial windows above V1 (n = 4 mice), PPC (n = 4), or fMC (n = 4; **Figure 3.3B**). We tested the necessity of each region during each of the three task epochs (stimulus, delay, response) with high temporal specificity (**Figure 3.3A**) by randomly interleaving 'laser ON' and 'laser OFF' trials in the same behavioral session. Photoinhibition of all three regions was capable of disrupting performance of the task, causing performance to drop to chance levels. However, the specific task epochs during which laser stimulation disrupted behavior differed for each brain region (**Figure 3.3C**).

As expected, V1 (**Figure 3.3C, D**, top row) was necessary during the stimulus epoch, with laser stimulation dramatically reducing the animal's ability to perform the task (stimulus: d'_{OFF} = 1.51 ± 0.11, d'_{ON}: 0.21 ± 0.11, $p < 10^{-3}$). However, inactivation of V1 during the delay or response epochs had no effect (delay: $\Delta d'$ = 0.47, p = 0.41, response: $\Delta d'$ = 0.23, p = 1.0). Activity in PPC (**Figure 3.3C, D**; middle row) was also necessary during the stimulus epoch (stimulus: d'_{OFF} = 1.81 ± 0.16, d'_{ON}: 0.27 ± 0.12, $p < 10^{-3}$). However, despite the high level of activity in many PPC neurons during the delay and response periods (see Figure 2.5), inactivation of PPC during these periods had no effect on behavior (delay: $\Delta d'$ = 0.03, p = 1.0, response: $\Delta d'$ = 0.24, p = 0.79). By contrast, suppression of fMC (**Figure 3.3C,D**; bottom row) during any of the three trial epochs completely abolished performance in the task (stimulus: d'_{OFF} = 1.37 ± 0.10, d'_{ON}: -0.05 ± 0.13, $p < 10^{-4}$; delay: d'_{OFF} = 1.60 ± 0.10, d'_{ON}: 0.32 ± 0.23, $p < 10^{-4}$; response: d'_{OFF} = 1.60 ± 0.17, d'_{ON}: -0.06 ± 0.05, $p < 10^{-3}$; paired t-test with Bonferroni correction used for all behavioral statistical tests).

One concern is that suppression of fMC could have disrupted performance in a trivial manner by preventing the execution of motor commands. This was indeed observed for inactivation of fMC during the response period, as licking behavior was abolished during laser ON trials (**Figure 3.3D**, bottom right). However, prevention of motor function cannot explain

decreased performance during stimulus or delay epoch inactivation, as light stimulation affected performance without decreasing lick rate (**Figure 3.3D**, bottom left & middle). The lick rate recovery after stimulus/delay epoch photoinhibition is somewhat puzzling given that the vast majority of the neurons in fMC are selective for Hit trials (see Figure 2.5B). In future experiments, it would be useful to image fMC activity after removal of photoinhibition to determine if the activity defaults to a "go" activity pattern after recovery.

Another concern is that the visual stimulus duration (2 s) is likely longer than necessary to perform the discrimination, meaning that the stimulus epoch could act as a mixed stimulus/delay epoch. To test this possibility, we trained mice on a similar task with a shortened stimulus duration of 250 ms (**Figure 3.4A**; performance declined to chance levels with shorter stimulus durations, data not shown). Stimulus period photoinhibition continues to disrupt behavioral performance for all three regions (**Figure 3.4B**), similar to the longer duration stimulus epochs **(Figure 3.3C**). Note that recovery of activity after photoinhibition can take several hundred milliseconds (**Figure 3.1B**; τ = 416 +/- 91 ms), so the results should be interpreted with caution. However, these data do demonstrate that fMC is necessary either during or very shortly after the visual stimulus is shown, highlighting the rapid conversion of sensory into motor information across cortical regions.

Our results thus show that all three regions are necessary during the stimulus epoch of the task. This is in line with our imaging results (see Figure 2.5) and population analyses (see Figure 2.6) in showing that task-relevant information reaches frontal motor regions relatively rapidly after stimulus presentation, and suggests that PPC may be involved with converting stimulus identity representations into behavioral choice representations early in the task (during the stimulus epoch), while fMC is necessary for maintaining and executing the behavioral choice (see Figure 2.7). Surprisingly, although there was robust task-related activity in PPC during the delay and response epochs, photoinhibition of PPC during these epochs had no impact on performance of the memory-guided task. Only fMC was necessary during the delay and response

epochs of the task, showing that fMC can act independently of V1 and PPC once the stimulus epoch processing is complete.

3.4.4 Photoinhibition of V1 and PPC during a 0 s delay task has similar effects

The apparent necessity of PPC in our task is intriguing, especially since recent work has brought the causal role of PPC in decision-making under contention. Although classical studies have implicated macaque PPC in perceptual decisions (Gold and Shadlen, 2007; Hanks et al., 2006; Shadlen and Newsome, 2001), recent muscimol experiments demonstrated that PPC can be inactivated with no effect on decision performance (Katz et al., 2016). Similarly, in rodents, PPC has been shown to be dispensable for decisions involving whisker-based object localization (Guo et al., 2014) or accumulation of auditory evidence (Erlich et al., 2015). However, PPC has been shown by other investigators to be causally involved in visual decisions during an accumulation task (Licata et al., 2016; Raposo et al., 2014), and a memory-guided navigational choice task (Harvey et al., 2012). Since inactivation of PPC in our visual delayed response task (see Figure 3.3C-D) appeared to disrupt behavior, we decided to investigate this further.

We used a simplified version of the visual discrimination task (Goard et al., 2016), with a 0 s delay between the stimulus and response period (**Figure 3.5A, B**). After presentation of the stimulus for 2 s, a retractable lick spout immediately became available to the animal for a 1.5 s response window. Licks to the target stimulus were rewarded with water (Hit trials), whereas licks to the non-target stimulus were punished with a small aversive drop of quinine (False Alarm trials). Mice (n = 4) achieved high levels of performance on the task (d-prime, 1.94 ± 0.12; mean ± SEM).

We again took advantage of an optogenetic approach to test the causal necessity of V1 and PPC in the task, using the transgenic *VGAT-EYFP-ChR2* mouse line, which express the light-sensitive opsin ChR2 in all inhibitory neurons (Zhao et al., 2011). Blue light stimulation (473nm)

of the cortex is sufficient to excite locally-projecting inhibitory neurons and thus suppress the activity of all nearby pyramidal cells (Guo et al., 2014), as shown in the example cell recorded *in vivo* in an anesthetized mouse (**Figure 3.5C**).

After training VGAT-EYFP-ChR2 mice on the behavioral task, we implanted bilateral windows over V1 (n = 3 mice) or over PPC (n = 4) and covered the rest of the skull with black cement, in order to restrict light stimulation to areas of interest (**Figure 3.5C**). We then tested the necessity of these two regions during a particular epoch of the task ("stimulus" or "response") by stimulating the cortex on randomly interleaved trials (**Figure 3.5D**). Photoinhibition of V1 during the stimulus period strongly disrupted behavioral performance during laser ON trials, by reducing responses during target stimuli to chance levels (**Figure 3.5D, top**; change in d', -1.68 ± 0.18). This result was expected given V1's known role in visual processing of orientation; both in our coarse discrimination task (Goard et al., 2016) and in detection paradigms (Glickfeld et al., 2013). By contrast, light stimulation of V1 during the response period had little effect. When we inactivated bilateral PPC, we found a similar pattern of results. Photoinhibition of PPC during the stimulus period also strongly disrupted behavioral performance during laser ON trials (**Figure 3.5D, bottom**; change in d', -1.37 ± 0.11). The disruption in performance could not be explained by changes in overall lick probability, as inactivation caused both an increase in false alarm rate and a decrease in hit rate (**Figure 3.5E**). By contrast, little effect was observed for stimulation of PPC during the response period. Examination of licking dynamics showed no clear effect of laser stimulation on the pattern of motor output for Hit or False Alarm trials (**Figure 3.5F**). This was true even when overall performance was reduced, as in the case of stimulus period inactivation.

These results confirm our prior report using a longer delay period (Goard et al., 2016) that inactivation of either V1 or PPC disrupts behavior only if applied during presentation of the visual stimulus (see Figure 3.3). Inactivation applied even immediately after stimulus offset has no effect on behavior, arguing strongly against a role for PPC in executing the motor choice.

3.4.5 Light stimulation of PPC does not activate putative inhibitory neurons in V1

The effects of manipulating V1 and PPC are remarkably similar. Both manipulations are specific to the stimulus period, and cause animals to lick indiscriminately at a rate intermediate between their baseline Hit and False Alarm rates. One legitimate concern, given the close proximity of PPC to V1 and other secondary visual areas (Garrett et al., 2014; Marshel et al., 2011; Wang and Burkhalter, 2007; Wang et al., 2012), is that this similarity could be due to nonspecific inactivation, or "inactivation spillover" (Brody and Hanks, 2016; Erlich et al., 2015). Other researchers have observed deficits on visual decision-making behavior when optically manipulating PPC (Licata et al., 2016), and they estimated the spatial extent of light stimulation using published calculators, concluding that that direct spread of light to V1 was unlikely. However, laser stimulation parameters can vary widely across studies and experimental preparations, and thus the spatial extent of light stimulation should be measured empirically.

Previous work, using the same mouse line and a similar stimulation paradigm but higher laser power densities, has shown that light stimulation can affect the activity of neurons as far as 1 or 2mm away (Guo et al., 2014). We previously measured the spatial extent of our stimulation approach using post-hoc immunohistochemistry and cFos labeling as a marker of neural activity (see Figure 3.1D). Light stimulation dramatically reduced cFos labeling under the cranial window, throughout all layers of the cortex, and extending less than 500μm from the edge of the window (Goard et al., 2016).

However, cFos is an indirect measure of neural activity that might not be sensitive to more subtle changes in firing rate. We therefore directly measured the effects of stimulating PPC on V1 spiking activity using single-unit electrophysiology (**Figure 3.6A, B**). Some single units (n = 7 across 3 mice) showed increased spiking when light stimulation was applied to V1 (average enhancement: 6.4-fold), presumably representing ChR2+ inhibitory cells that are directly activated by light (**Figure 3.6D, E**). Consistent with this assumption, these cells tended to exhibit narrow

spike widths (**Figure 3.6C**; 5 of 7 units with spike widths less than 0.3 ms), an extracellular signature commonly used to distinguish excitatory and inhibitory neurons (McCormick et al., 1985; Mitchell et al., 2007; Runyan and Sur, 2013; Swadlow, 2003). Light stimulation of PPC had no effect on the majority of these cells (5 of 7 cells), with one cell showing significant enhancement and one showing significant suppression. This demonstrates that photoinhibition of PPC is unlikely to cause direct light activation of inhibitory cells in V1.

3.4.6 Light stimulation of PPC causes moderate suppression of V1 neurons

In addition to the V1-enhanced cells, we recorded many units (n = 34 across 6 mice) that were strongly suppressed by direct light stimulation of V1 (**Figure 3.7A,B**; 93.4% suppression). These V1-suppressed neurons had broader spike widths on average than the V1-enhanced population (**Figure 3.6C**). When light stimulation was applied to PPC, we found that many of these cells were moderately suppressed (average suppression: 60.7%; 25 of 34 cells with significant suppression, $p < 0.01$, Wilcoxon rank-sum test). We hypothesize that this moderate suppression is not due to direct light activation of nearby inhibitory neurons in V1, but rather to the loss of long-range excitatory input from PPC. Consistent with this hypothesis, the normalized spike rate (stimulation divided by control) was not significantly correlated with distance (**Figure 3.8B**; $r = 0.13$, $p = 0.48$), as would be expected for a polysynaptic effect as opposed to a direct effect of light-mediated inactivation. Additionally, normalized spike rate was weakly but not significantly correlated with recording depth ($r = 0.26$, $p = 0.13$).

Photoinhibition of PPC therefore has an indirect and moderate effect on V1 activity. This is not unexpected given the interconnected nature of PPC and V1 (Wang et al., 2012), although it does warrant caution in interpretation of our behavioral results, as with any acute manipulation experiment including pharmacological manipulations or acute lesions (Otchy et al., 2015). Disruption of behavior during PPC inactivation could be due to disruption of computations within

PPC, but it could be also explained by loss of feedback excitation onto V1. Either way, PPC appears to provide signals necessary for performance of the task, whether by performing sensorimotor computations or by providing feedback excitation onto V1. Both V1 and PPC are necessary exclusively during the stimulus period and not during execution of the motor response.

3.5 Discussion

Sensorimotor decisions involve both transformation and maintenance of information, and it remains a major goal of systems neuroscience to understand how neural circuits perform these functions. As a first step towards this goal, we manipulated the activity of neurons in multiple cortical regions during a memory-guided sensorimotor decision task to discover which areas were fundamental to each of these processes.

Our first major finding is that the frontal motor cortex, and not posterior parietal cortex, appears to be responsible for maintenance of the motor plan during the delay period, adding to several other recent studies (Erlich et al., 2011; Guo et al., 2014; Hernandez et al., 2010; Li et al., 2015). Indeed, while photoinhibition of fMC during the delay period results in complete disruption of behavioral performance, photoinhibition of PPC during the delay period has no significant effect (**Figure 3.3**). This suggests that, at least in motor planning tasks, PPC may be more important for mapping stimulus identity to behavioral choice than for maintaining the choice in short-term memory stores. However, additional research will be necessary to determine if the PPC is indeed the locus for stimulus-response mapping in these tasks. Another intriguing finding is that the dependency on fMC arises quite early in the task (**Figure 3.3**), within hundreds of milliseconds from stimulus onset (**Figure 3.4**). This result appears to conflict with earlier findings using unilateral inactivation of motor cortex (Guo et al., 2014; Hanks et al., 2015), but is consistent with a recent study finding that delay-period motor cortical activity can be recovered if the contralateral hemisphere is spared (Li et al., 2016). It is unclear whether these differences should be attributed

to differences in behavioral task or stereotactic coordinates. Our chosen coordinates lie slightly posterior and medial to the motor area described by one group to be important for licking behaviors in mice (Guo et al., 2014; Komiyama et al., 2010), but the size of our stimulation window likely overlaps with this area. Nonetheless, the early dependency on fMC activity suggests a very rapid conversion of stimulus information into a motor plan in our task. An alternative possibility is that disruption of fMC affects other areas during the stimulus period that are necessary for performance. While we did not measure activity in V1 or PPC during fMC inactivation, the lack of strong direct connectivity of fMC with either area (Zingg et al., 2014) makes the possibility of off-target effects less likely.

Second, we find that inactivation of PPC during the stimulus epoch of the task disrupts behavior, whether using a memory-guided version of the task (**Figure 3.3**), a very brief stimulus presentation (**Figure 3.4**), or a 0 s delay between stimulus and choice (**Figure 3.5**). This adds to prior work which indicate a role for rodent PPC in other visual decision tasks (Harvey et al., 2012; Licata et al., 2016; Raposo et al., 2014) by specifying an early involvement of PPC in such tasks. But our results conflict with perturbation studies in rodents (Erlich et al., 2015; Guo et al., 2014) and primates (Katz et al., 2016), which found little role for PPC in other decision tasks. To resolve these discrepancies, one must first ask the question: what exactly is being disrupted when we shine light in PPC during the stimulus period?

One possibility that our electrophysiological recordings rule out is that of direct light spread to V1. Putative inhibitory neurons in V1, which exhibit narrow spike waveforms and are activated by direct stimulation of V1, are not excited by light stimulation of PPC (**Figure 3.6**). These measurements are bolstered by post-hoc immunohistochemical measurements of cFos, which show reduced labeling across all layers in a spatial region that extends at most 500 μm laterally from the cranial window (**Figure 3.1**). However, it must be acknowledged that the precise definition of rodent PPC, its homology to primate cortex, and its anatomical relationship to

extrastriate visual areas, remains a matter of debate. It is possible that rodent PPC is homologous to primate PPC in being unnecessary for visual decisions (Katz et al., 2016). In this case, behavioral deficits in rodents may arise only in visual tasks because of nonspecific inactivation of separate, nearby extrastriate areas (Brody and Hanks, 2016). An alternative possibility is that rodent PPC shares features with primate extrastriate cortex, and thus deficits on visual decisions observed by us and others could be due to direct disruption of visual or visuomotor computations within PPC (Licata et al., 2016). Distinguishing these possibilities will require improved resolution and delineation of these areas, perhaps using more precise retinotopic mapping (Garrett et al., 2014) or noise correlation analyses (Kiani et al., 2015).

Still another possibility is that PPC inactivation induces off-target effects in other areas, particularly V1. Indeed, our electrophysiological recordings show that the role of PPC cannot be easily assessed independently from that of V1. The two brain areas are directly and indirectly interconnected (Wang et al., 2012), and manipulation of PPC activity does moderately affect V1 spiking as well (**Figure 3.7**). This warrants caution in interpretation of our behavioral results, as with any acute manipulation including optical or pharmacological manipulations (Otchy et al., 2015). However, the reduction in V1 activity was far from complete, with an average ~40% of visually-evoked spikes spared. It is unlikely that an incomplete disruption of V1 activity can fully explain the large behavioral deficit that we observed when inactivating PPC. Low contrast visual stimuli can produce even weaker neural responses, and even so our mice are able to perform visual discrimination at such low contrasts (see Figure 4.4B). The lack of distance dependence (**Figure 3.8**) and effect on putative inhibitory cells further argues that such reduction is a network-level effect.

We therefore argue that PPC is causally involved in the task, whether by directly computing visuomotor transformations, or by indirectly providing permissive feedback excitation to V1. Fully distinguishing whether PPC plays an instructive versus a merely permissive role will

require additional experiments. Specifically, one should selectively lesion PPC, and then allow time for homeostatic recovery of activity levels in V1 and other areas that receive input from PPC. If task performance is preserved after lesion (without any task-specific practice), then one can conclude that PPC is not in fact required for task performance, but provides permissive input to V1 or other areas that are essential (Otchy et al., 2015; Panzeri et al., 2017). A third and intriguing possibility is that PPC and V1 may jointly participate in an inter-areal network that processes and transforms stimulus information in a manner that is necessary for task performance. Such a perspective, while speculative, goes beyond attributing specific functions to individual areas, and can be tested in future work using targeted measurement and manipulation of specific projection pathways between areas (Chen et al., 2013a; Li et al., 2015; Znamenskiy and Zador, 2013).

Although the necessity of PPC activity during the stimulus period requires future inquiry, we observed a clear negative result when inactivating PPC in the delay and response periods. Does this activity play any functional role, or is it merely epiphenomenal? Delay-period PPC activity could reflect feedback inputs from motor and frontal regions that drive the behavior, or it may point to a role for PPC in learning (Law and Gold, 2008). Another intriguing alternative comes from unilateral disruption experiments in both primates (Katz et al., 2016; Wardak et al., 2002; Wilke et al., 2012) and rodents (Erlich et al., 2015), which demonstrate clear biases in free choice behavior even though the same manipulations have little effect on instructed sensory trials. Activity in PPC could therefore be more important for biasing attention or behavior in less demanding task conditions, as seen also in humans with hemi-spatial neglect (Kerkhoff, 2001).

One caveat is that disruption of behavioral performance in a go/no-go task may be hard to interpret due to the asymmetry of the task design. We provided controls to demonstrate that deficits in performance were not due to distraction by the light stimulus itself, prolonged temporal effects, or non-specific changes in overall lick rate (except for fMC during the response period).

However, we cannot rule out the possibility that behavioral disruption is due to laser-induced changes in motivation, arousal, or perceptual criterion.

An important next step is to determine the role of V1, PPC, and fMC in a task that requires memory of the stimulus, such as a delayed match-to-sample task (Liu et al., 2014; Miller et al., 1996; Romo et al., 1999). It seems probable that in such tasks, PPC or perhaps even primary sensory cortex would play a greater role during delay-period maintenance. Given the importance of both sensory and motor short-term memory, both delayed-response and delayed-comparison tasks warrant further study.

Finally, this study focused on cortical regions located on the dorsal surface of the mouse brain, as they are readily accessible to optogenetic inhibition. However, there are likely to be additional cortical and subcortical regions involved in this task, including prefrontal cortex (Funahashi et al., 1989; Fuster and Alexander, 1971; Kojima and Goldman-Rakic, 1982), thalamus (Fuster and Alexander, 1973), superior colliculus (Kopec et al., 2015), and basal ganglia (Ding and Gold, 2012; Kawagoe et al., 1998). These regions could be further investigated with a similar approach as invasive (Andermann et al., 2013; Barretto and Schnitzer, 2012; Zorzos et al., 2012) and noninvasive (Filonov et al., 2011; Mittmann et al., 2011; Prakash et al., 2012) techniques for optical interrogation of deeper structures become available.

3.6 Contributions and Acknowledgements

M.J.G. and G.N.P. designed experiments with input from M.S.; M.J.G., G.N.P., and J.W. performed the behavioral training and photoinhibition experiments; G.N.P. developed the behavior software; M.J.G. performed the histology experiments; G.N.P performed the electrophysiology experiments; M.J.G. and G.N.P. analyzed the data with input from M.S.; M.J.G., G.N.P., and M.S. wrote the manuscript.

We thank J. Sharma, S. Ramirez, B. Crawford, A. Boesch, V. Li, C. Le, and T. Emery for technical assistance; S. El-Boustani and R. Huda for comments on the manuscript; H. Robertson, Q. Chen and G. Feng for providing VGAT-ChR2-EYFP mice. This work was supported by the NIH (M.J.G., F32-EY023523 and K99-MH104259; M.S., R01-EY007023 and U01-NS090473); the NSF (G.N.P., Graduate Research Fellowship; M.S., EF1451125); the Simons Center for the Social Brain (M.S.), and the Picower Institute Innovation Fund (M.J.G.; M.S.).

3.7 References

Andermann, M.L., Gilfoy, N.B., Goldey, G.J., Sachdev, R.N., Wolfel, M., McCormick, D.A., Reid, R.C., and Levene, M.J. (2013). Chronic cellular imaging of entire cortical columns in awake mice using microprisms. Neuron *80*, 900-913.

Andermann, M.L., Kerlin, A.M., and Reid, R.C. (2010). Chronic cellular imaging of mouse visual cortex during operant behavior and passive viewing. Frontiers in cellular neuroscience *4*, 3.

Andersen, R.A., and Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits. Neuron *63*, 568-583.

Barretto, R.P., and Schnitzer, M.J. (2012). In vivo optical microendoscopy for imaging cells lying deep within live tissue. Cold Spring Harb Protoc *2012*, 1029-1034.

Bauer, R.H., and Fuster, J.M. (1976). Delayed-matching and delayed-response deficit from cooling dorsolateral prefrontal cortex in monkeys. J Comp Physiol Psychol *90*, 293-302.

Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. J Neurosci *31*, 913-921.

Bisley, J.W., Zaksas, D., and Pasternak, T. (2001). Microstimulation of cortical area MT affects performance on a visual working memory task. J Neurophysiol *85*, 187-196.

Brody, C.D., and Hanks, T.D. (2016). Neural underpinnings of the evidence accumulator. Current opinion in neurobiology *37*, 149-157.

Chafee, M.V., and Goldman-Rakic, P.S. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. J Neurophysiol *79*, 2919-2940.

Chafee, M.V., and Goldman-Rakic, P.S. (2000). Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. J Neurophysiol *83*, 1550-1566.

Chen, J.L., Carta, S., Soldado-Magraner, J., Schneider, B.L., and Helmchen, F. (2013a). Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. Nature *499*, 336-340.

Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V.*, et al.* (2013b). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295-300.

Constantinidis, C., and Steinmetz, M.A. (1996). Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. J Neurophysiol *76*, 1352-1355.

de Lafuente, V., and Romo, R. (2005). Neuronal correlates of subjective sensory experience. Nature neuroscience *8*, 1698-1703.

di Pellegrino, G., and Wise, S.P. (1993). Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. J Neurosci *13*, 1227-1243.

Ding, L., and Gold, J.I. (2010). Caudate encodes multiple computations for perceptual decisions. The Journal of neuroscience : the official journal of the Society for Neuroscience *30*, 15747-15759.

Ding, L., and Gold, J.I. (2012). Separate, causal roles of the caudate in saccadic choice and execution in a perceptual decision task. Neuron *75*, 865-874.

Erlich, Jeffrey C., Bialek, M., and Brody, Carlos D. (2011). A Cortical Substrate for Memory-Guided Orienting in the Rat. Neuron *72*, 330-343.

Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. Elife *4*.

Filonov, G.S., Piatkevich, K.D., Ting, L.M., Zhang, J., Kim, K., and Verkhusha, V.V. (2011). Bright and stable near-infrared fluorescent protein for in vivo imaging. Nat Biotechnol *29*, 757-761.

Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. Nature *443*, 85-88.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J Neurophysiol *61*, 331-349.

Fuster, J.M., and Alexander, G.E. (1971). Neuron activity related to short-term memory. Science *173*, 652-654.

Fuster, J.M., and Alexander, G.E. (1973). Firing changes in cells of the nucleus medialis dorsalis associated with delayed response behavior. Brain Res *61*, 79-91.

Garrett, M.E., Nauhaus, I., Marshel, J.H., and Callaway, E.M. (2014). Topography and areal organization of mouse visual cortex. J Neurosci *34*, 12587-12600.

Gisquet-Verrier, P., and Delatour, B. (2006). The role of the rat prelimbic/infralimbic cortex in working memory: not involved in the short-term maintenance but in monitoring and processing functions. Neuroscience *141*, 585-596.

Glickfeld, L.L., Histed, M.H., and Maunsell, J.H. (2013). Mouse primary visual cortex is used to detect both orientation and contrast changes. J Neurosci *33*, 19416-19422.

Goard, M.J., Pho, G.N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife *5*, 558.523.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annual review of neuroscience *30*, 535-574.

Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014). Flow of cortical activity underlying a tactile decision in mice. Neuron *81*, 179-194.

Hanks, T.D., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nat Neurosci *9*, 682-689.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220-223.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62-68.

Hernandez, A., Nacher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vazquez, Y., Camarillo, L., and Romo, R. (2010). Decoding a perceptual decision process across cortex. Neuron *66*, 300-314.

Horwitz, G.D., and Newsome, W.T. (1999). Separate signals for target selection and movement specification in the superior colliculus. Science *284*, 1158-1161.

Kampa, B.M., Roth, M.M., Gobel, W., and Helmchen, F. (2011). Representation of visual scenes by local neuronal populations in layer 2/3 of mouse visual cortex. Front Neural Circuits *5*, 18.

Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. Nature *535*, 285-288.

Kawagoe, R., Takikawa, Y., and Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. Nat Neurosci *1*, 411-416.

Kerkhoff, G. (2001). Spatial hemineglect in humans. Prog Neurobiol *63*, 1-27.

Kiani, R., Cueva, C.J., Reppas, J.B., Peixoto, D., Ryu, S.I., and Newsome, W.T. (2015). Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. Neuron *85*, 1359-1373.

Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat Neurosci *2*, 176-185.

Kojima, S., and Goldman-Rakic, P.S. (1982). Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. Brain Res *248*, 43-49.

Komiyama, T., Sato, T.R., O'Connor, D.H., Zhang, Y.X., Huber, D., Hooks, B.M., Gabitto, M., and Svoboda, K. (2010). Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. Nature *464*, 1182-1186.

Kopec, C.D., Erlich, J.C., Brunton, B.W., Deisseroth, K., and Brody, C.D. (2015). Cortical and Subcortical Contributions to Short-Term Memory for Orienting Movements. Neuron *88*, 367-377.

Law, C.T., and Gold, J.I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat Neurosci *11*, 505-513.

Li, C.S., Mazzoni, P., and Andersen, R.A. (1999). Effect of reversible inactivation of macaque lateral intraparietal area on visual and memory saccades. J Neurophysiol *81*, 1827-1838.

Li, N., Chen, T.W., Guo, Z.V., Gerfen, C.R., and Svoboda, K. (2015). A motor cortex circuit for motor planning and movement. Nature *519*, 51-56.

Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. Nature *532*, 459-464.

Licata, A.M., Kaufman, M.T., Raposo, D., Ryan, M.B., Sheppard, J.P., and Churchland, A.K. (2016). Posterior parietal cortex guides visual decisions in rats. bioRxiv.

Liu, D., Gu, X., Zhu, J., Zhang, X., Han, Z., Yan, W., Cheng, Q., Hao, J., Fan, H., Hou, R.*, et al.* (2014). Medial prefrontal activity during delay period contributes to learning of a working memory task. Science *346*, 458-463.

Marshel, J.H., Garrett, M.E., Nauhaus, I., and Callaway, E.M. (2011). Functional specialization of seven mouse visual cortical areas. Neuron *72*, 1040-1054.

McCormick, D.A., Connors, B.W., Lighthall, J.W., and Prince, D.A. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. J Neurophysiol *54*, 782-806.

McNaughton, B.L., Mizumori, S.J., Barnes, C.A., Leonard, B.J., Marquis, M., and Green, E.J. (1994). Cortical representation of motion during unrestrained spatial navigation in the rat. Cereb Cortex *4*, 27-39.

Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. J Neurosci *16*, 5154-5167.

Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. Neuron *55*, 131-141.

Mittmann, W., Wallace, D.J., Czubayko, U., Herb, J.T., Schaefer, A.T., Looger, L.L., Denk, W., and Kerr, J.N. (2011). Two-photon calcium imaging of evoked activity from L5 somatosensory neurons in vivo. Nat Neurosci *14*, 1089-1093.

Murakami, M., Vicente, M.I., Costa, G.M., and Mainen, Z.F. (2014). Neural antecedents of self-initiated actions in secondary motor cortex. Nat Neurosci *17*, 1574-1582.

Nakamura, K., and Colby, C.L. (2000). Visual, saccade-related, and cognitive activation of single neurons in monkey extrastriate area V3A. J Neurophysiol *84*, 677-692.

Nitz, D.A. (2006). Tracking route progression in the posterior parietal cortex. Neuron *49*, 747-756.

Oh, S.W., Harris, J.A., Ng, L., Winslow, B., Cain, N., Mihalas, S., Wang, Q., Lau, C., Kuan, L., Henry, A.M.*, et al.* (2014). A mesoscale connectome of the mouse brain. Nature *508*, 207-214.

Otchy, T.M., Wolff, S.B., Rhee, J.Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S.M., and Olveczky, B.P. (2015). Acute off-target effects of neural circuit manipulations. Nature *528*, 358-363.

Panzeri, S., Harvey, C.D., Piasini, E., Latham, P.E., and Fellin, T. (2017). Cracking the Neural Code for Sensory Perception by Combining Statistics, Intervention, and Behavior. Neuron *93*, 491-507.

Prakash, R., Yizhar, O., Grewe, B., Ramakrishnan, C., Wang, N., Goshen, I., Packer, A.M., Peterka, D.S., Yuste, R., Schnitzer, M.J.*, et al.* (2012). Two-photon optogenetic toolbox for fast inhibition, excitation and bistable modulation. Nature methods *9*, 1171-1179.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat Neurosci *17*, 1784-1792.

Romo, R., Brody, C.D., Hernandez, A., and Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. Nature *399*, 470-473.

Runyan, C.A., and Sur, M. (2013). Response Selectivity Is Correlated to Dendritic Structure in Parvalbumin-Expressing Inhibitory Neurons in Visual Cortex. Journal of Neuroscience *33*, 11724-11733.

Sachidhanandam, S., Sreenivasan, V., Kyriakatos, A., Kremer, Y., and Petersen, C.C. (2013). Membrane potential correlates of sensory perception in mouse barrel cortex. Nat Neurosci *16*, 1671-1677.

Sakurai, Y., and Sugimoto, S. (1985). Effects of lesions of prefrontal cortex and dorsomedial thalamus on delayed go/no-go alternation in rats. Behavioural brain research *17*, 213-219.

Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. Journal of neurophysiology *86*, 1916-1936.

Siegel, M., Buschman, T.J., and Miller, E.K. (2015). Cortical information flow during flexible sensorimotor decisions. Science *348*, 1352-1355.

Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Coding of intention in the posterior parietal cortex. Nature *386*, 167-170.

Sreenivasan, K.K., Curtis, C.E., and D'Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. Trends Cogn Sci *18*, 82-89.

Super, H., Spekreijse, H., and Lamme, V.A. (2001). A neural correlate of working memory in the monkey primary visual cortex. Science *293*, 120-124.

Swadlow, H.A. (2003). Fast-spike interneurons and feedforward inhibition in awake sensory neocortex. Cereb Cortex *13*, 25-32.

Wang, Q., and Burkhalter, A. (2007). Area map of mouse visual cortex. The Journal of comparative neurology *502*, 339-357.

Wang, Q., Sporns, O., and Burkhalter, A. (2012). Network analysis of corticocortical connections reveals ventral and dorsal processing streams in mouse visual cortex. J Neurosci *32*, 4386-4399.

Wardak, C., Olivier, E., and Duhamel, J.R. (2002). Saccadic target selection deficits after lateral intraparietal area inactivation in monkeys. J Neurosci *22*, 9877-9884.

Whitlock, J.R., Sutherland, R.J., Witter, M.P., Moser, M.B., and Moser, E.I. (2008). Navigating from hippocampus to parietal cortex. Proceedings of the National Academy of Sciences of the United States of America *105*, 14755-14762.

Wilke, M., Kagan, I., and Andersen, R.A. (2012). Functional imaging reveals rapid reorganization of cortical activity after parietal inactivation in monkeys. Proceedings of the National Academy of Sciences of the United States of America *109*, 8274-8279.

Wilson, N.R., Runyan, C.a., Wang, F.L., and Sur, M. (2012). Division and subtraction by distinct cortical inhibitory networks in vivo. Nature, 1-6.

Zagha, E., Xinxin, G., and McCormick, D.A. (2015). Competing Neural Ensembles in Motor Cortex Gate Goal-Directed Motor Output. Neuron *88*, 565-577.

Zhao, S., Ting, J.T., Atallah, H.E., Qiu, L., Tan, J., Gloss, B., Augustine, G.J., Deisseroth, K., Luo, M., Graybiel, A.M.*, et al.* (2011). Cell type-specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. Nature methods *8*, 745-752.

Zingg, B., Hintiryan, H., Gou, L., Song, M.Y., Bay, M., Bienkowski, M.S., Foster, N.N., Yamashita, S., Bowman, I., Toga, A.W.*, et al.* (2014). Neural networks of the mouse neocortex. Cell *156*, 1096-1111.

Znamenskiy, P., and Zador, A.M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. Nature *497*, 482-485.

Zorzos, A.N., Scholvin, J., Boyden, E.S., and Fonstad, C.G. (2012). Three-dimensional multiwaveguide probe array for light delivery to distributed brain circuits. Opt Lett *37*, 4841-4843.

3.8 Figures

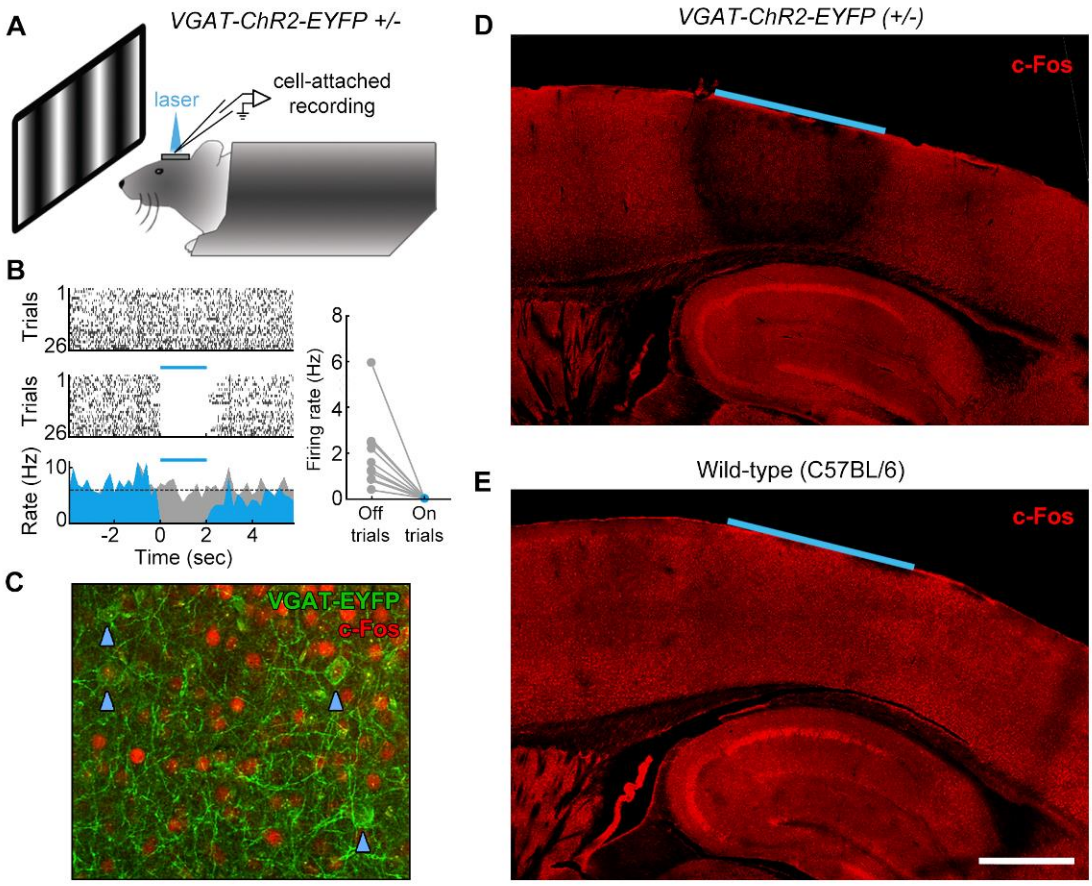

Figure 3.1. Characterization of photoinhibition in VGAT-ChR2-EYFP mice.

(**A**) Schematic of cell-attached recording set-up to test effect of blue light on regular spiking neurons in VGAT-ChR2-EYFP mice. (**B**) Response of example layer 5 regular spiking neuron with high baseline firing rate on interleaved laser off (top) and laser on (bottom) trials. Blue light completely silences neural activity during all trials (bottom). Complete silencing was observed for all regular spiking neurons recorded (right, *n* = 10 neurons). (**C**) Cortical slice from VGAT-ChR2-EYFP mouse showing constitutive c-Fos expression (red) in putative excitatory (EYFP-negative) neural somata. Note that c-Fos expresses at low levels in inhibitory interneurons (EYFP-positive; blue arrowheads). (**D**) Laser stimulation in VGAT-ChR2-EYFP mice dramatically reduces c-Fos expression (red) in a local region underneath the cranial window. Reduction of c-Fos spreads farthest in middle layers, but is generally limited to a few hundred microns from the window edge. Reduction of c-Fos was not observed in subcortical regions. The light blue line indicates window location (2mm diameter window). (**E**) Identical laser stimulation in wild-type mice does not affect c-Fos expression (red) underneath the cranial window. The light blue line indicates window location (2mm diameter window). Scale bar, 1 mm.

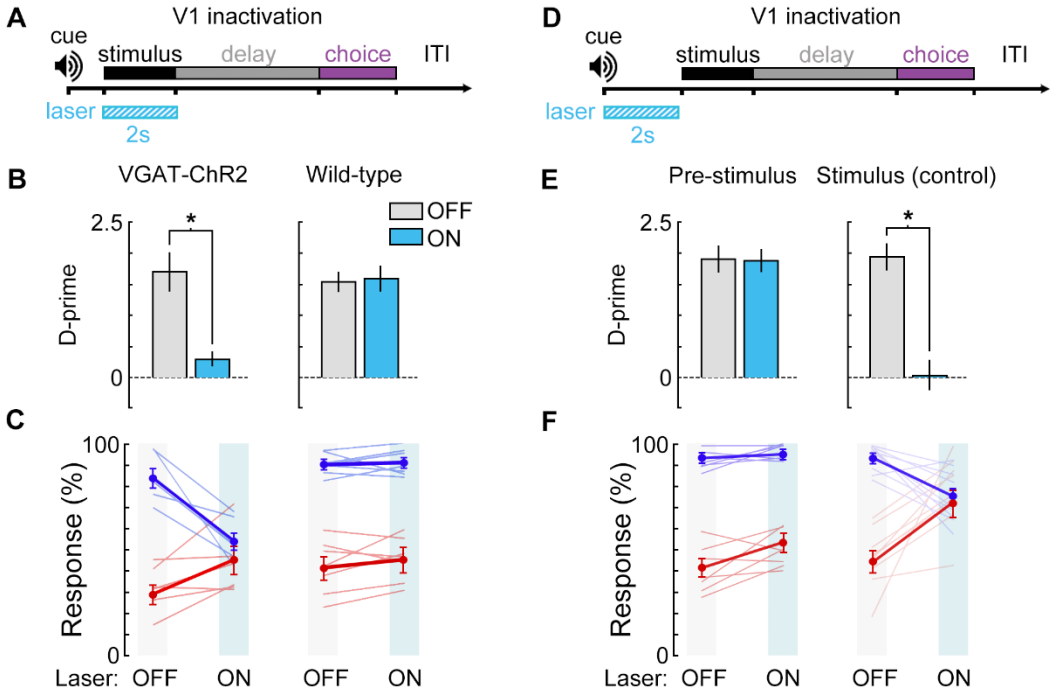

Figure 3.2. Photoinhibition control experiments.

(**A**) Bilateral V1 laser stimulation in VGAT-ChR2-EYFP and control (wild-type) mice during the stimulus epoch of a 0 s delay version of the task. (**B,C**) Bilateral laser stimulation significantly impairs behavior in a VGAT-ChR2-EYFP mouse, causing performance to fall near chance levels ($\Delta d'$ = -1.41, p = 0.002, n = 6 sessions). Laser stimulation has no effect on performance for a wild-type control mouse ($\Delta d'$ = +0.14, p = 0.27, n = 8 sessions). (**D**) To ensure that laser stimulation does not exert effects on performance lasting beyond the laser stimulation period, we investigated the effect of laser stimulation during a 2 s period immediately prior to visual stimulus onset in a VGAT-ChR2-EYFP mouse. (**E,F**) Laser stimulation during the 2 s prior to the stimulus epoch had no significant effect on performance ($\Delta d'$ = -0.15, p = 0.51, n = 8 sessions) compared to laser stimulation during the stimulus period for the same mouse ($\Delta d'$ = -1.93, p = 0.003, n = 8 sessions). Paired t-tests used for all statistical tests.

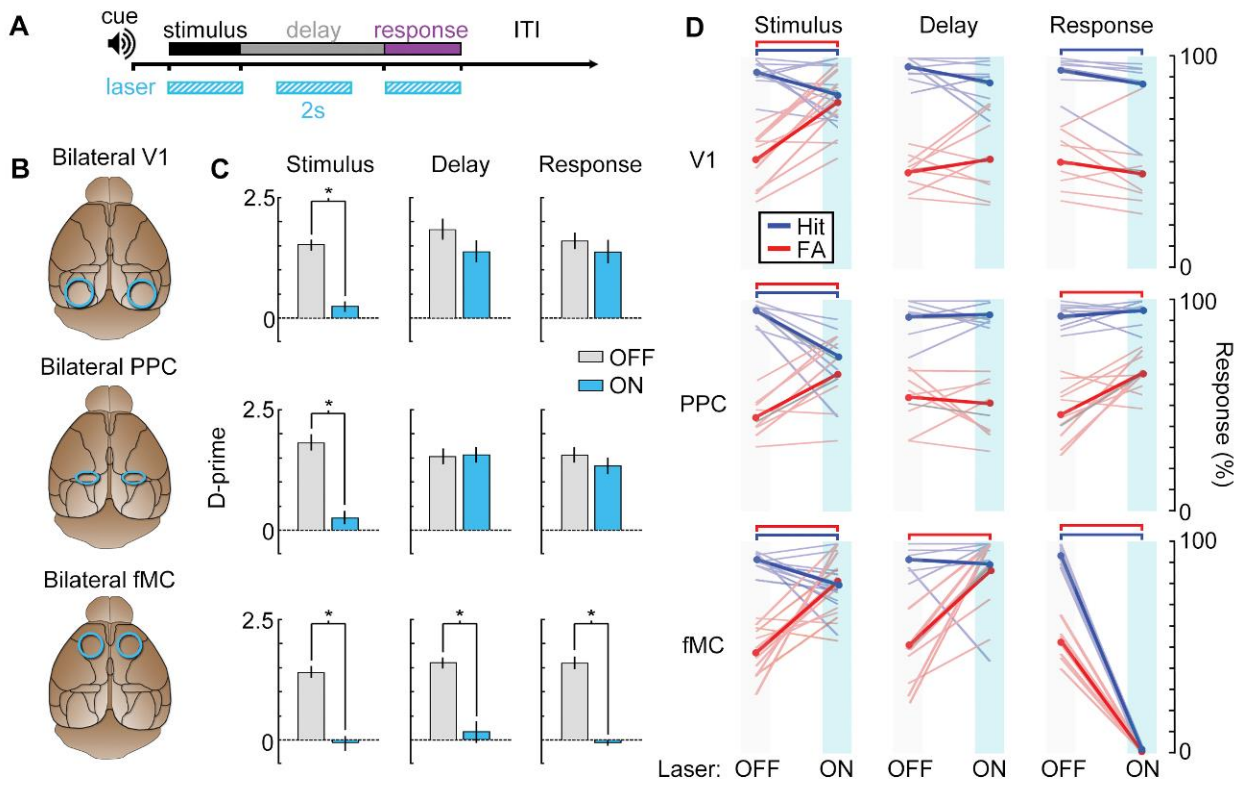

Figure 3.3. Photoinhibition of specific cortical regions during task performance.

(**A**) Continuous blue light stimulation (473nm, 2s) was applied on interleaved trials during either the stimulus, delay, or response epochs of the task. (**B**) Glass windows were implanted bilaterally over V1 (top), PPC (middle) or fMC (bottom) of VGAT-ChR2-YFP mice. Superficial blue light stimulation silences activity in nearby pyramidal cells (Figure 3.1), effectively silencing the exposed region. (**C**) Behavioral performance (d-prime) during interleaved laser OFF and ON trials for each brain region and trial epoch. Photoinhibition of V1 (top row) or PPC (middle row) significantly disrupts behavioral performance only when applied during the stimulus period (left column), whereas photoinhibition of fMC (bottom row) during any epoch of the task significantly decreases d-prime (*, $p < 0.05$, *t*-test with Bonferroni correction). (**D**) Effect of laser stimulation on hit rate (blue) and false alarm rate (red). Lightly colored lines represent individual behavioral sessions. Colored bars at top indicate statistical significance for each group ($p < 0.05$, *t*-test with Bonferroni correction). For all plots, error bars indicate ± S.E.M.

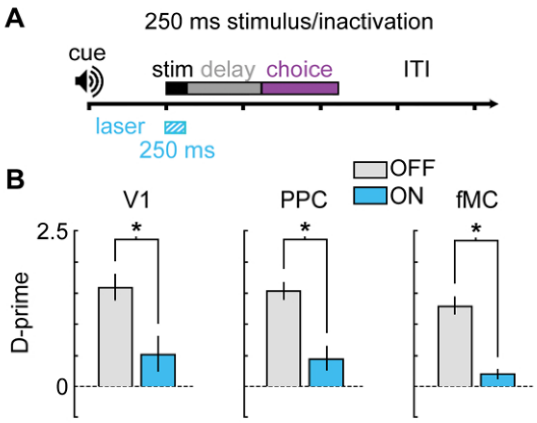

Figure 3.4. Bilateral photoinactivation during brief (250 ms) visual stimuli.

(**A**) Mice were trained on a version of the task with a brief (250 ms) visual stimulus, followed by a 1 sec delay. Pre-stimulus period (1 s), response period (1.5 s), and ITI (3 s) were unchanged from the standard task. Laser was on for 250 ms during the stimulus. (**B**) Performance was significantly impaired for V1 ($\Delta d'$ = -1.08, p = 0.048, n = 5 sessions), PPC ($\Delta d'$ = -1.09, p = 0.009, n = 6 sessions), and fMC ($\Delta d'$ = -1.08, p = 0.003, n = 5 sessions).

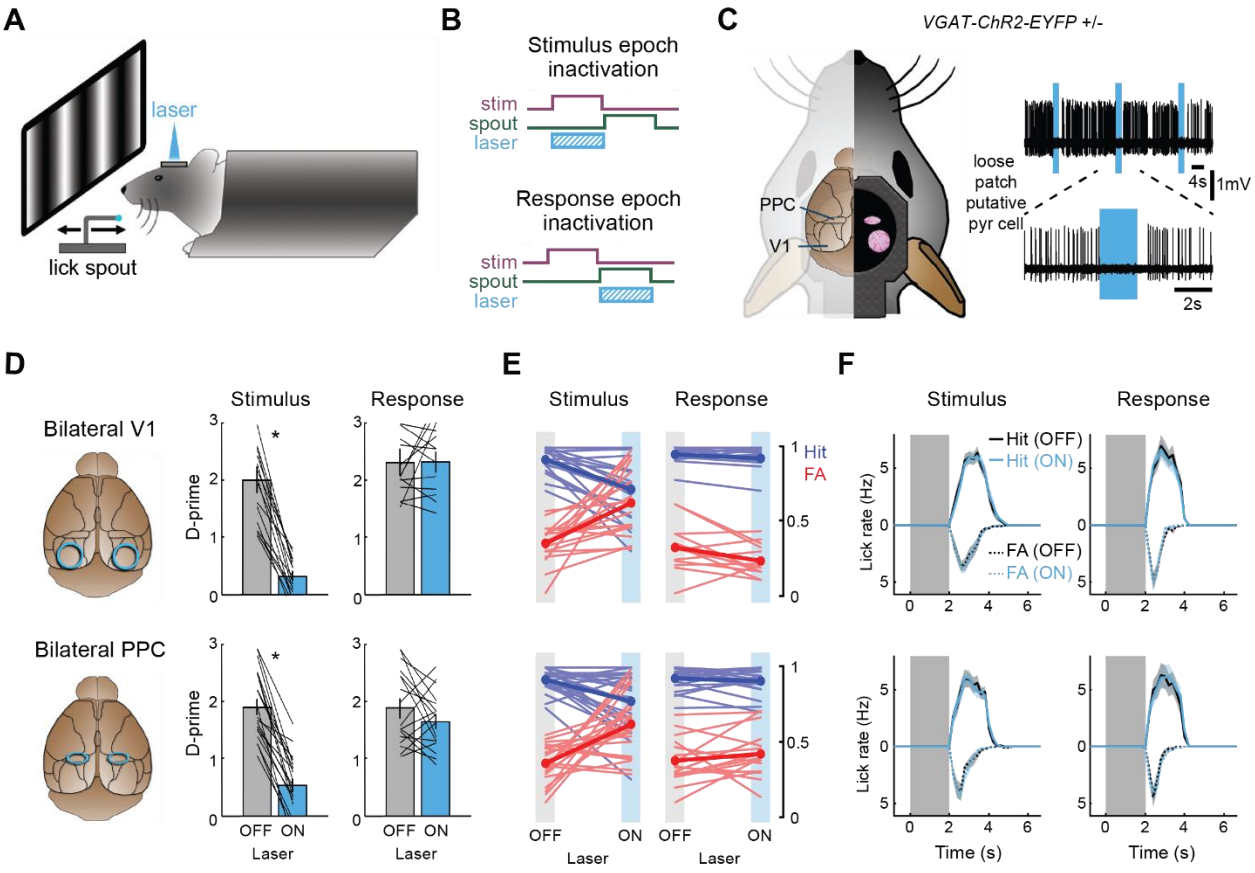

Figure 3.5. Photoinhibition of V1 and PPC during a 0 s delay task has similar effects.

(**A**) Experimental setup for photoinhibition in head-fixed mice performing a visual discrimination task. A retractable lick spout was used to restrict lick responses to a specific epoch of the task. (**B**) Trial structure of the 0-second delay task. After a brief auditory preparatory cue, a drifting grating was presented for 2 s. The retractable lick spout was presented immediately after the stimulus turned off for at least 1.5 s. Continuous blue light stimulation (473nm, 2s) was applied on interleaved trials during either the stimulus or response epochs of the task. (**C**) Schematic illustrating inactivation method (left). Glass windows cut to size were implanted over bilateral PPC and/or V1. Example loose patch recording demonstrating complete suppression of a putative pyramidal neuron by light (right). Photostimulation of GABAergic neurons expressing ChR2 in VGAT-ChR2-EYFP mice can be used to transiently inactivate local pyramidal cells. (**D**) Behavioral performance (d-prime) during interleaved laser Off and On trials for both task epochs (columns) and brain regions (rows). *, $p < 0.001$, Wilcoxon signed-rank test. Lines indicate individual sessions. (**E**) Effect of photoinhibition on lick rates for both task epochs and brain regions. Inactivation of bilateral V1 (top row, n = 3 mice) disrupts behavior by reducing hit rates (blue) and increasing false alarm rates (red), but only when applied during the stimulus epoch (left) and not during the response epoch (right). Inactivation of bilateral PPC (bottom row, n = 4 mice) similarly causes a reduction in hit rates and increase in false alarm rates only during the stimulus epoch (left) and not during the response epoch (right). Lines indicate individual sessions. (**F**) Effect of photoinhibition on licking dynamics for both task epochs (columns) and brain regions (rows). Average lick peri-stimulus time histograms with (blue) and without (black) laser stimulation for both Hit trials (solid) and False Alarm trials (dotted). Lick PSTHs on False Alarm trials are plotted with an inverted Y-axis for illustration. Shading indicates SEM across sessions. Photoinhibition had no effect on licking dynamics for any stimulation condition.

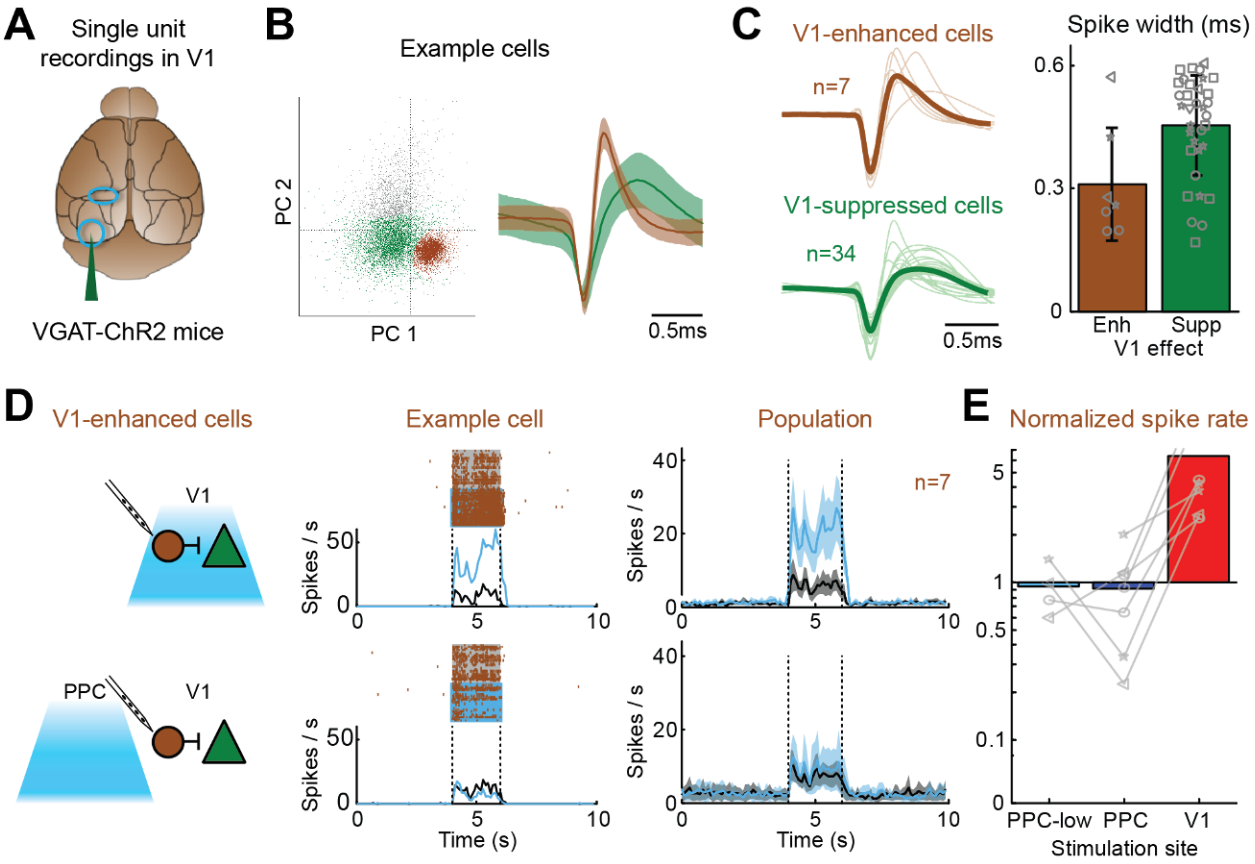

Figure 3.6. Light stimulation of PPC does not activate putative inhibitory neurons in V1.

(**A**) Single-unit recordings (using a 16-channel silicon linear probe) were made in V1 of untrained VGAT-ChR2-EYFP mice under isoflurane anesthesia during laser stimulation of either V1 or PPC. (**B**) Example regular-spiking (green) and fast-spiking (brown) units recorded from the same channel. Left, principal components analysis on detected waveforms reveals clearly separable clusters. Right, average waveforms have different spike widths. Shading indicates STD. (**C**) V1-enhanced cells are more likely to be fast-spiking than V1-suppressed cells. Left, waveforms of all V1-enhanced (top, brown, n = 7) and V1-suppressed cells (bottom, green, n = 34), with overlaid population averages. Right, spike widths of V1-enhanced and V1-suppressed cells. Bars denote mean ± SEM. Markers and shapes indicate individual units from different mice (n=6). (**D**) Effect of V1 stimulation (top row) or PPC stimulation (bottom row) on V1-enhanced cells. Left column, trial-sorted spike raster and peri-stimulus time histogram for the example fast-spiking cell in (B), with (blue line) and without (black line) laser stimulation. Grey shading or dashed lines indicate visual stimulation (2 s). Blue shading indicates laser stimulation (2.2 s, onset precedes visual stimulation by 0.1 s). Right column, average peri-stimulus time histograms across all V1-enhanced cells (n=7). Shading indicates SEM. V1-enhanced cells do not show consistent enhancement by PPC stimulation. (**E**) Spike rate of V1-enhanced cells, normalized to no-laser trials for V1 (red) and PPC (dark blue) stimulation at full power (34 mW). PPC was also stimulated at low power (PPC-low, light blue, 19mW) for a subset of cells (n=4). Bars denote mean. Markers and shapes indicate individual units from different mice. Normalized spike rates below 0.1 were rounded down to 0. PPC stimulation had an inconsistent effect, with only 2 of 7 cells showing significant enhancement (n=1) or suppression (n=1).

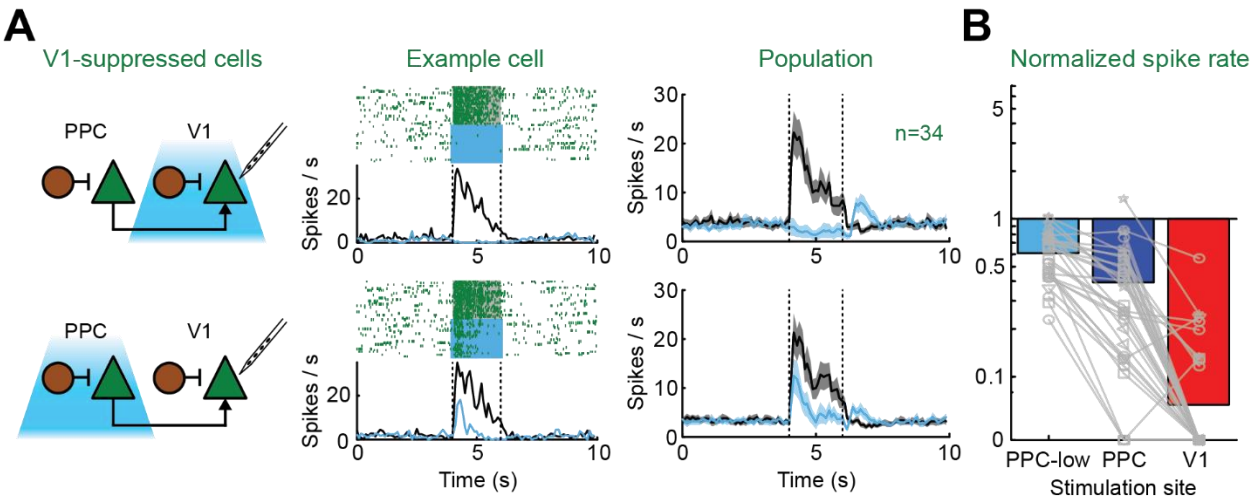

Figure 3.7. Light stimulation of PPC causes moderate suppression of V1 neurons.

(**A**) Effect of V1 stimulation (top row) or PPC stimulation (bottom row) on V1-suppressed cells. Left column, trial-sorted spike raster and peri-stimulus time histogram for an example regular-spiking cell, with (blue line) and without (black line) laser stimulation. Grey shading or dashed lines indicate visual stimulation (2 s). Blue shading indicates laser stimulation (2.2 s, onset precedes visual stimulation by 0.1 s). Right column, average peri-stimulus time histograms across all V1-suppressed cells (n=34). Shading indicates SEM. V1-suppressed cells are moderately suppressed by PPC stimulation. (**B**) Spike rate of V1-suppressed cells, normalized to no-laser trials for V1 (red) and PPC (dark blue) stimulation at full power (34 mW). PPC was also stimulated at low power (PPC-low, light blue, 19mW) for a subset of cells (n=28). Bars denote mean. Markers and shapes indicate individual units from different mice. Normalized spike rates below 0.1 were rounded down to 0. V1 stimulation always had a greater effect (93.4% suppression) than PPC stimulation at high power (60.7%) or low power (39.7%).

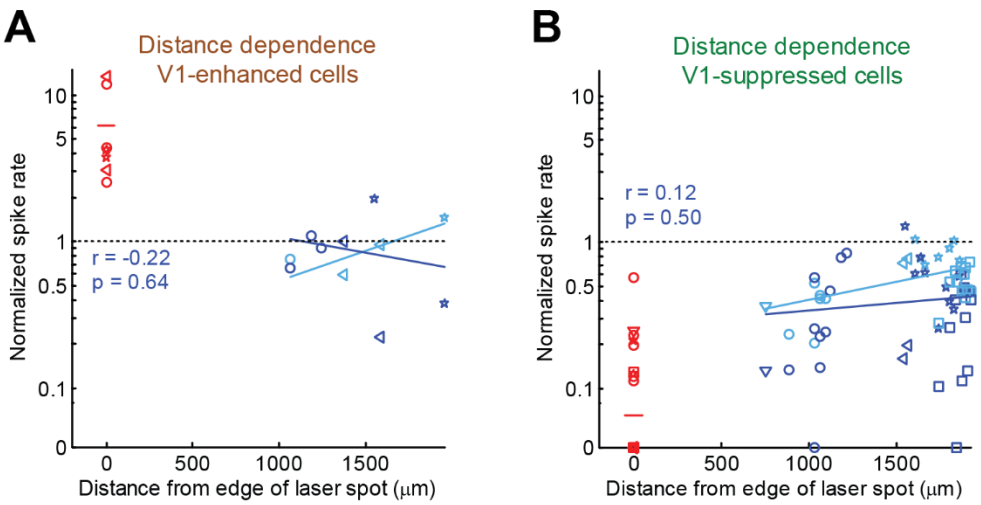

Figure 3.8. Lack of distance-dependence suggests indirect, network-level suppression of V1.

Normalized spike rate versus distance for V1-enhanced (**A**) and V1-suppressed cells (**B**) for PPC stimulation at both high power (dark blue) and low power (light blue), and for V1 stimulation (red). Distances from edge of laser spot were estimated from probe insertion location, location of PPC window, and location of electrode site along the probe. Red line indicates average effect of V1 stimulation. Blue lines indicate best-fit linear regression for high power (dark blue) and low power (light blue) stimulation of PPC. There was no significant correlation ($p = 0.50$) between normalized spike rate (at high power) and distance.

Chapter 4: Task-dependent representations of stimulus and choice in mouse parietal cortex.

4.1 Summary

The posterior parietal cortex (PPC) has been implicated in perceptual decisions, but whether its role is specific to sensory processing or sensorimotor transformation is not well understood. To distinguish these possibilities, we trained mice to perform a visual discrimination task and imaged the activity of PPC populations during both engaged behavior and passive viewing. Unlike neurons in primary visual cortex (V1), which responded robustly to stimuli in both conditions, most neurons in PPC responded exclusively during task engagement. However, PPC responses were heterogeneous, with a smaller subset of neurons exhibiting stimulus-driven, contrast-dependent responses in both conditions. To test whether PPC responses primarily encoded the stimulus or the learned sensorimotor contingency, we imaged the same neurons before and after re-training mice on a reversed task contingency. Unlike V1 neurons, most PPC neurons exhibited a dramatic shift in selectivity after re-training and reflected the new sensorimotor contingency, while a smaller subset of neurons preserved their stimulus selectivity. Mouse PPC is therefore strongly task-dependent, contains heterogeneous populations sensitive to stimulus and choice, and may play an important role in the flexible transformation of sensory inputs into motor commands.

4.2 Introduction

Perceptual decision-making involves multiple cognitive processes, including processing of sensory stimuli, accumulation of evidence, and transformation of sensory information into an appropriate motor plan. Although many brain regions have been implicated in perceptual decisions, dissociating their individual contribution to these different processes remains a challenge.

In particular, the posterior parietal cortex (PPC) has been hypothesized to play a key role in at least some types of decision tasks in both primates (Gold and Shadlen, 2007; Shadlen and Newsome, 2001) and in rodents (Goard et al., 2016; Harvey et al., 2012; Raposo et al., 2014). However, the specific role of rodent parietal cortex, and whether its function is homologous to that of primates, remains unclear. Recent electrophysiological recordings in rats during auditory decisions (Hanks et al., 2015) recapitulated findings in primates that PPC neurons can reflect the accumulated evidence for a decision. However, pharmacological inactivation of PPC during the same auditory task has no effect on behavioral performance (Erlich et al., 2015), and a minimal role of PPC was also found for whisker-based decisions in mice (Guo et al., 2014b). These findings are in contrast to studies involving visual decision tasks, where activity in PPC not only encodes information about the decision (Goard et al., 2016; Morcos and Harvey, 2016; Raposo et al., 2014) but is also causally necessary for behavior (Goard et al., 2016; Harvey et al., 2012; Licata et al., 2016).

The apparent specificity of behavioral deficits to the visual modality has led some to argue that rodent PPC may ultimately be more homologous to extrastriate cortex in processing sensory signals that are accumulated elsewhere for decision-making (Licata et al., 2016). Indeed, both anatomical projection studies (Wang and Burkhalter, 2007; Wang et al., 2012), as well as functional mapping studies (Garrett et al., 2014; Marshel et al., 2011) indicate that PPC may overlap with or contain a group of retinotopically-organized extrastriate areas that are rostral to V1. An alternative possibility is that PPC may play a specific role in the mapping of visual stimuli to motor commands. If this were the case, one may expect that activity in PPC would highly task-dependent, and may even flexibly re-map depending on learned sensorimotor contingencies (Freedman and Assad, 2006).

To distinguish these possibilities, we used population calcium imaging to measure activity in PPC and in the primary visual cortex (V1) during a go/no-go lick-based visual discrimination task. Having previously demonstrated the necessity of PPC in the performance of this task (Goard

et al., 2016), we sought in this work to investigate its specific role in either sensory processing or sensorimotor transformation by manipulating task engagement and learned task contingencies. V1 neurons had robust visual responses during both task engagement and passive viewing of stimuli that remained stable after task contingency reversal. By contrast, PPC responses were largely specific to task performance and were highly sensitive to task contingency. Our results are consistent with a role of the mouse posterior parietal cortex in transforming sensory information to motor commands during perceptual decisions.

4.3 Experimental Procedures

4.3.1 Surgical procedures

All experiments were carried out in mice of either sex using protocols approved by Massachusetts Institute of Technology's Animal Care and Use Committee and conformed to National Institutes of Health guidelines. Data were collected from adult (3-5 months old) wild-type (C57BL/6; n = 15) mice. The animals were housed on a 12 hour light/dark cycle in cages of up to 5 animals before the implants, and individually after the implants. All surgeries were conducted under isoflurane anesthesia (3.5% induction, 1.5-2.5% maintenance). Meloxicam (1 mg kg^{-1}, subcutaneous) was administered pre-operatively and every 24 hours for 3 days to reduce inflammation. Surgical procedures for the first headplate implant surgery are identical to those described in Section 2.3.1. After head plate implantation, mice recovered for at least five days before beginning water restriction.

After behavioral training was complete, animals were taken off water restriction for five days before undergoing a second surgery to implant the imaging window. Procedures for anesthetic administration and post-operative care were identical to the first surgery. The dental acrylic and silicon elastomer covering the targeted region were removed using a drill burr. The

skull surface was then cleaned and a craniotomy was performed over left V1/PPC, leaving the dura intact. Neurons were labeled with a genetically-encoded calcium indicator by microinjection (Stoelting) of 50 nl AAV2/1.Syn.GCaMP6s.WPRE.SV40 (University of Pennsylvania Vector Core; diluted to a titer of 10^{12} genomes ml^{-1}) 300 µm below the pial surface. Between two and five injections were made in each exposed region, centered at V1 (4.2 mm posterior, 2.5 mm lateral to Bregma) and PPC (2 mm posterior, 1.7 mm lateral to Bregma). Since the viral expression spreads laterally from the injection site, exact stereotaxic locations were photographed through the surgical microscope for determining imaging areas. Finally, a cranial window was implanted over the craniotomy and sealed first with silicon elastomer then with dental acrylic. The cranial windows were made of two rounded pieces of coverglass (Warner Instruments) bonded with optical glue (NOA 61, Norland). The bottom piece was circular or oval, custom cut according to cortical region(s) (V1: 2.5 mm x 2.5 mm; PPC: 1 mm x 2 mm; V1 + PPC: 4 mm x 4 mm, anterior-posterior x medial-lateral) and fit snugly in the craniotomy. The top piece was a larger circular coverglass (3-5 mm, depending on size of bottom piece) and was bonded to the skull using dental acrylic. Mice recovered for five days before commencing water restriction.

4.3.2 Behavioral tasks

We trained mice to perform a head-fixed go/no-go visual discrimination task. Stimuli consisted of full-contrast sine wave gratings (spatial frequency: 0.05 cycles deg^{-1}; temporal frequency: 2 Hz) drifting at either 0° (upwards, target, Stimulus A) or 90° (rightwards, non-target, Stimulus B) away from vertical. Stimuli were presented to the right eye alone by placing the screen at an oblique angle to the animal. Behavioral training and testing was implemented with custom software written in Matlab (Mathworks) using Psychtoolbox-3 (Kleiner et al., 2007) and Data Acquisition toolbox. Spout position was controlled by mounting the spout apparatus on a pneumatically-driven sliding linear actuator (Festo) controlled by two solenoids. Licks were detected using an infrared

emitter/receiver pair (Digikey) mounted on either side of the retractable lick spout. Mice were water-restricted and earned most of their daily ration (1mL) during training.

Detailed information about the behavioral training procedures is provided in Section 4.3.3. An auditory cue tone (5 kHz, 0.5 s, 65 dB SPL) indicated the beginning of each trial. After a 1 s delay, a visual stimulus was presented for 2 s. At the end of the stimulus epoch, the spout was rapidly moved within reach of the tongue, and remained within reach for 1.5 s. Correct licks during this period were rewarded with 5-8 µl water and a brief reward tone (10 kHz, 0.1 s). Licks to the non-target stimulus were punished with a white noise auditory stimulus alone (early training) or white noise plus 1-3 µl of 5mM quinine hydrochloride in water (late training). This concentration was chosen to deter licking to non-targets without causing mice to lose motivation altogether. At the end of the response epoch, the spout was then rapidly retracted and remained out of reach until the next trial (3 s inter-trial interval).

During imaging experiments, blocks of engaged behavior trials were alternated with blocks of passive viewing. Blocks were 5-10 min in duration (40-80 trials per block). During passive blocks, the spout was out of reach for the duration of the block. A few extra passive trials were given (without imaging) before each passive block to ensure that mice did not expect spout presentation during all imaged passive trials. The sequence of target and non-target stimuli presented for a given passive block was matched to the sequence of stimuli used for the preceding engaged block. In some cases, instead of alternating between engaged and passive blocks, the engaged blocks were all grouped together at the beginning of the session, followed by an equal number of consecutive passive blocks. No difference in results was observed for alternating versus grouped blocks.

Some mice (n = 6) were trained on a variable contrast version of the task. On each trial, the stimulus was randomly set to one of six contrasts (2, 4, 8, 16, 32, or 64%), regardless of whether the stimulus was a target or non-target. The mouse therefore could not predict the contrast of the stimulus from trial to trial.

Some mice (n = 3) were re-trained after initial imaging experiments on a reversed reward contingency. Reward contingency was switched abruptly, with reward given for licks to Stimulus B and no reward for licks to Stimulus A. Because mice were quickly discouraged by the reversal, no quinine punishment was initially given. Additionally, the reward tone (10 kHz, 0.1 s) was paired with the onset of the new target stimulus (Stimulus B) early in re-training, in order to encourage licking. Three of the five mice trained on this reversed contingency achieved criterion performance after re-training for 10 ± 2 days.

For all tasks, behavioral d-prime was computed by norminv(Hit rate) – norminv(False alarm rate), where norminv() is the inverse of the cumulative normal function (Carandini and Churchland, 2013; Green and Swets, 1966). Values of Hit and False alarm rate were truncated between 0.01 and 0.99, setting the maximum d' to 4.65.

4.3.3 Behavioral training

Mice were trained once a day, 5-6 days per week. Mice were trained in successive stages, with advancement to the next stage contingent on correct performance (see **Supplementary Materials**): 1) Mice receiving reward any time they licked the spout. 2) Trial structure was initiated by having an auditory cue tone, followed by a visual stimulus (100% targets), followed by an inter-trial interval. Mice were only rewarded for licks during the visual stimulus. 3) Once mice exhibited preferential licking during the stimulus, the target rate was reduced over several sessions from 100% to 50%. At this point, the non-target was a static grating orientated orthogonally to the target. Licks during non-targets were punished with white noise or white noise plus quinine. 4) Once mice exhibited the ability to discriminate target and non-target gratings ($d' > 1$ and $R_{HIT} - R_{FA} > 30\%$ for several sessions, where R_{HIT} and R_{FA} are the hit and false alarm rate, respectively), the temporal frequency of the non-target grating was increased. 5) Spout withdrawal was introduced. At first the spout was extended within range before the stimulus appeared, then spout extend time

was gradually delayed until after the stimulus had turned off. 6) Some mice were also trained on the variable contrast version of the task. Stimuli of lower contrast were gradually added and randomly interleaved in with higher contrast stimuli, until mice could achieve significant performance even at the lowest contrast (2%). Mice that failed to fully learn the task within 150 sessions or showed signs of infection were removed from the study. Some mice were additionally trained on a delayed response version of the task and re-used in another study (Goard et al., 2016) which added to the training time.

Once mice reached high levels of performance at the final stage of the task ($d' > 1.5$ and $R_{HIT} - R_{FA} > 50\%$), they were removed from water restriction for window implantation. Mice reached criterion performance after an average of 92 ± 11 sessions. After recovery from window implantation surgery, they were re-trained to a level of high performance (2-7 days) before beginning experimental sessions. Any sessions with poor performance were discarded (minimum performance criterion: $d' > 1$ and $R_{HIT} - R_{FA} > 30\%$).

4.3.4 Two-photon imaging

Procedures for two-photon calcium imaging were identical to those described in Section 2.3.3. Briefly, GCaMP6s fluorescence was imaged 14-35 days after virus injection using Prairie Ultima IV 2-photon microscopy system with a resonant galvo scanning module (Bruker). We used a 16x/0.8 NA microscope objective (Nikon), which was mounted on a Z-piezo (Bruker) for volume scanning. Four 441 x 512 pixel imaging planes separated by 20 or 25 μm were imaged sequentially at a stack rate of 5 Hz for 5-10 min imaging sessions. Laser power ranged from 40-75 mW at the sample depending on GCaMP6s expression levels. During imaging experiments, the polypropylene tube supporting the mouse was suspended from the behavior platform with high tension springs (Small Parts) to dampen movement.

4.3.5 Image analysis

Calcium imaging data were acquired using PrairieView acquisition software and sorted into multi-page TIF files. All analyses were performed using custom scripts written either in ImageJ or MATLAB (Mathworks).

Images were first corrected for X-Y movement by registration to a reference image (the pixel-wise mean of all frames) using 2-dimensional cross correlation. To identify responsive neural somata, a pixel-wise activity map was calculated as previously described (Ahrens et al., 2012). Neuron cell bodies were identified using local adaptive threshold and iterative segmentation. Automatically-defined ROIs were then manually checked for proper segmentation in a MATLAB-based graphical user interface (allowing comparison to raw fluorescence and activity map images). To subtract the influence of local neuropil on somatic signals, the fluorescence in the somata was estimated as $F_{corrected\_soma}(t) = F_{raw\_soma}(t) - 0.7 \times F_{neuropil}(t)$, where $F_{neuropil}$ was the defined as the fluorescence in the region 0-15 mm from the ROI border (excluding other ROIs) (Chen et al., 2013). $\Delta F/F$ for each neuron was calculated as $\Delta F/F_t = (F_t - F_0)/ F_0$, with F_0 defined as the mode of the raw fluorescence density distribution.

To align ROIs between different imaging sessions across days (**Figure 4.7**, **Figure 4.8**), we used a semi-automated method similar to prior work (Huber et al., 2012). First, for each plane, anchor points were manually defined by visual comparison of the two average projection images. These anchor points helped to define a predicted displacement vector field that would be used to map coordinates from one session to the other. For each coordinate, the predicted vector was defined by the average (weighted inversely by distance) of the vectors for all defined anchor points.

Next, for each ROI, a square region (~4x the size of the ROI) around the ROI was selected. To determine the displacement across sessions, we computed the normalized cross-correlation

of this square with the average projection of the other session. This was multiplied point-by-point with a mask that decayed gradually with distance from the predicted displacement vector, and then smoothed with a 2-D Gaussian filter. The peak of the resulting image was taken to be the actual displacement vector of the ROI. This process biases the displacement of each ROI towards the vector predicted from the manually-defined anchor points. Finally, any ROIs with a computed displacement vector that differed by greater than 5 pixels from the predicted vector were flagged for manual inspection, and then either redrawn or removed.

4.3.6 Analysis of task-driven responses

After image preprocessing and $\Delta F/F$ extraction, traces were sorted by trial type (hit, miss, correct reject, false alarm) and condition (engaged, passive). The baseline response (1 s before stimulus onset) was subtracted from each trial. A neuron was considered task responsive if its mean $\Delta F/F$ during the last 1.6 s (8 frames) of the stimulus period was significantly ($p < 0.01$, t-test) greater than the pre-stimulus baseline (1 s), for either hit or correct reject trials. Neurons also had to meet a signal-to-noise criterion, needing a trial-averaged response that exceeded a threshold of at least 2 standard deviations above baseline during either the stimulus or choice period. All further analyses are based on responses during the last 1.6 s of the stimulus period, unless noted otherwise.

In some sessions, an unnoticed bug in the behavior software caused the non-target stimulus to be presented during both the stimulus and choice period (3.5 s total). Because most of our analyses were confined to the 2 s stimulus period, the bug should not dramatically alter the results. Additionally, within-mouse comparisons of behavior showed no significant difference in performance with or without the bug. Nonetheless, we re-analyzed the data excluding these sessions and found no major differences in our key findings in **Figure 4.2** and **Figure 4.3**. We

therefore included these sessions in our stimulus-period analyses, but not in plots or analyses that involve the choice period (**Figure 4.2B**, **D**; **Figure 4.3**).

Neurons were marked as target- or non-target-preferring (**Figure 4.2E**) based on their mean response during Engaged trials. Neurons were marked as task-gated if they did not exhibit a significant response to their preferred stimulus during Passive trials (**Figure 4.2F**).

All comparative indices (engagement modulation index, contrast modulation index, selectivity index) were computed using a receiver operating characteristic (ROC) analysis, which quantifies the ability of an ideal observer to discriminate between trial types based on single trial responses (Britten et al., 1996; Green and Swets, 1966). Each index was derived from the area under the ROC curve (AUC), and defined as $2 \times (\mathrm{AUC} - 0.5)$; this value ranged from -1 to 1 (Raposo et al., 2014). Engagement modulation index (**Figure 4.2H**, **Figure 4.4D, Figure 4.8D**) was computed by comparing the stimulus period response (last 1.6 s) during engaged and passive conditions. A value of +1 indicates that a cell always fired more on Engaged trials. Contrast modulation index (**Figure 4.4C**) was computed by comparing responses at the highest and lowest contrasts, with positive values indicating preference for high contrasts. Selectivity index (**Figure 4.2G**) was computed by comparing target and non-target responses, with positive values indicating preference for target stimuli. Finally, selectivity index was computed separately before and after reversal training (**Figure 4.7E-H**) by comparing responses to the Stimulus A (red, original target) and Stimulus B (blue, original non-target), with positive values indicating preference for the Stimulus A.

To determine whether these comparative indices were significant (for either individual neurons or for the whole population), we used a permutation test. We shuffled the labels for each trial and recomputed the index 2000 times to create a distribution of indices that could have arisen by chance. Indices outside the center 95% interval of this distribution were considered significant.

4.3.7 Error trial analysis

Data from error trials are included only in **Figure 4.3**, **Figure 4.6**, and **Figure 4.8C**, and were taken from behavioral sessions in which the mouse committed at least five misses or five false alarms. Because mice had a bias towards licking, the number of Miss trials was small, and we focused our analysis on False Alarm trials (except in **Figure 4.3E-F**). Error modulation index (**Figure 4.8C**) was computed in the same manner as other comparative indices (see Section 2.3.5) but by comparing stimulus period responses on False Alarm trials versus Correct Reject trials, with positive values indicating preference for False Alarm trials.

Selectivity for stimulus and choice (**Figure 4.3E-H**, **Figure 4.6C-F**) were computed using an ROC-based analysis across different time bins (Bennur and Gold, 2011; Hernandez et al., 2010). For stimulus encoding, Hit trials were compared with False Alarm trials. For choice encoding, False Alarm trials were compared with Correct Reject trials. The auROC was computed by comparing the two distributions of responses at each time point (200 ms bins) during the trial. A permutation test was used to determine the percentage of neurons with significant selectivity ($p < 0.05$), and bootstrapping across sessions was used to determine confidence intervals on these percentages.

4.3.8 Contrast task analysis

For data acquired during the variable contrast task, neurons were considered significantly responsive if the mean $\Delta F/F$ during the stimulus period was significantly above threshold for at least two of the six contrasts of the same stimulus. We focused our analyses on target-preferring neurons, which were included if their mean response across contrasts was greater for Hit (target) trials compared to Correct Reject trials.

Single neuron contrast response functions were fit to the hyperbolic ratio function (**Figure 4.4E**), also known as the Naka-Rushton function (Albrecht and Hamilton, 1982):

$$R(C) = R_{max}\frac{C^n}{C^n + {C_{50}}^n} + R_0$$

where $R(C)$ is the neural response as a function of contrast, R_{max} is the saturation point, C_{50} is the contrast at the half-saturation point, R_0 is the baseline response, and n is an exponent that determines the steepness of the curve. The responses for both Engaged and Passive conditions were fit simultaneously, with n constrained to be constant across conditions, by minimizing the sum (across data points) of the squared error between the model and the data, divided by the variance of that data point.

We evaluated the goodness of fit for each neuron using a bootstrap hypothesis test, as detailed by others (Carandini et al., 1997). Briefly, we tested the null hypothesis that the mean of the probability distribution underlying the neural responses was identical to the predictions of the model. We measured the observed prediction error (e_{obs}) and computed the probability of observing an error at least as large if the null hypothesis were true. To sample from a distribution that conformed to the null hypothesis, we shifted the data such that the mean responses equaled the model predictions, and drew 1000 bootstrap samples from this dataset, computing a prediction error (e_i) for each. The proportion of samples for which the prediction error was larger than e_{obs} is the achieved significance level. For neurons with an achieved significance level below 10% ($p<0.1$), there was sufficiently strong evidence against the model, and therefore these neurons were excluded from further analysis.

Contrast modulation index was computed to compare responses on high (64%, 32%) versus low (2%, 4%) contrasts. A permutation test was used to assess significance ($p<0.05$) by shuffling trial labels 2000 times, and comparing the measured index to the shuffled distribution of indices.

4.3.9 Reverse contingency task analysis

We imaged from the same field of neurons before and after reversal of reward contingency. The two sessions were separated by an average time interval of 16 ± 1 days. A semi-automated method was used to align ROIs between the two sessions (see Section 4.3.5). Neurons were included for analysis only if a significant response ($p < 0.01$) to either stimulus was observed both before and after reversal.

4.3.10 General statistics

Unless otherwise noted, all measures are reported as mean ± SEM. Due to very large sample sizes, very small p-values ($<10^{-9}$) were approximated as $p < 10^{-9}$ as a lower bound on reasonable probabilities.

4.4 Results

4.4.1 Imaging calcium responses in V1 and PPC during engaged task performance and passive viewing

We trained mice on a head-fixed lick/no-lick visual discrimination task (**Figure 4.1A, B**), similar to previous designs (Andermann et al., 2010; Goard et al., 2016; Pinto et al., 2013). Water-restricted mice discriminated between a target stimulus (horizontal grating drifting upwards, 0° from vertical, Stimulus A) which was rewarded with water, and a non-target stimulus of an orthogonal orientation (vertical drifting upwards rightwards, 90°, Stimulus B). Lick responses to the non-target grating were discouraged by punishment with a small, aversive drop of quinine. A retractable lick spout was presented immediately after stimulus presentation (2 s), and retracted on every trial. This restricted the animal's lick response to the "response" period (1.5 s), and allowed us to

separately assess perception and action (**Figure 4.1B**). Video recording of the mice during the stimulus period confirmed that mice withheld licking until the spout became present during the response epoch (data not shown). Mice (n = 15) achieved high levels of performance on the task (**Figure 4.1C**; d-prime, 2.23 ± 0.18; mean ± SEM), with a bias towards licking, resulting in more false alarms than misses.

We have previously demonstrated using a version of this task with a 4 s delay period (Goard et al., 2016), that inactivation of either V1 or PPC during the stimulus period disrupts behavioral performance, whereas inactivation during the response period has no effect (see Figure 3.3). Additional inactivation experiments using the same 0 s delay task presented here confirmed this result (see Figure 3.5). Stimulus-period activity in PPC is therefore necessary for the task, but it is unclear whether such activity is important for sensory processing, decision formation, or action selection. In order to distinguish between these possibilities, we measured neural activity in PPC and V1 under two different conditions: during engagement in the behavioral task as well as during passive visual stimulation (**Figure 4.1B**). We reasoned that neurons important for sensory processing would show a robust response to stimuli during both passive and engaged conditions, whereas decision- or action-encoding neurons would show strong modulation by behavioral state.

We used two-photon microscopy and a volumetric imaging approach to image hundreds of neurons simultaneously in either V1 or PPC (see Section 4.3.4). After completion of behavioral training, we injected AAV2/1 syn-GCaMP6s (Chen et al., 2013b) into the two areas under stereotaxic guidance. We used a resonant scanning system combined with a z-piezo to concurrently record activity from several hundreds of GCaMP6-infected cells within layer 2/3 in a volume comprising four planes (850μm x 850μm) 20μm apart in depth, at an overall stack rate of 5 Hz. Images were corrected for X-Y movement, and fluorescence traces were extracted from semi-automatically generated ROIs based on a pixel-wise activity map (see Section 4.3.5).

To investigate the effects of behavioral performance on neural responses, we imaged the same neurons in alternating blocks (5-10 min, 40-80 trials) of engaged behavior and passive viewing (**Figure 4.1B**). During “engaged” behavior trials, the spout was extended on each trial during the response period. During “passive” trials, the spout was not extended during the response period but remained withdrawn and inaccessible from the animal. To avoid extraneous stimulus confounds, no additional cue was provided to signal engaged or passive blocks. Nonetheless, mice rapidly became aware after the first 1-2 trials of a block whether the spout would be available for a behavioral response, as confirmed in video recordings by the complete lack of licking during passive blocks (data not shown). Discrimination performance was similarly high (d’, 2.22 ± 0.18) for both the first and subsequent “engaged” blocks.

We imaged from an average of 606 cells (range: 257-1057) in each population within either V1 or PPC, of which an average 135 cells (range: 21-364) exhibited significant task-related responses. Pilot experiments using transgenic mice with tdTomato expressed in PV+ and SOM+ interneurons revealed that calcium signals from these neurons were generally too weak to be measured with volume scanning, and therefore the vast majority of task-responsive cells were likely to be excitatory pyramidal neurons. We first analyzed only the neural responses measured during correct Engaged trials, or during Passive viewing. In V1, many neurons showed a robust and reliable response to either the target or non-target stimulus, though some were modulated by engagement in the task. For example, one target-preferring cell showed a reliable passive response to the target stimulus that was moderately enhanced during task performance (**Figure 4.1D**, left). A non-target preferring cell, however, exhibited a suppressed response during engagement in the task (**Figure 4.1D**, right). By contrast, PPC neurons exhibited much stronger responses during task engagement compared to passive viewing, specifically for target stimuli. For example, one target-preferring cell had robust activity only during engagement (**Figure 4.1E**, left), whereas another cell showed relatively weak passive responses that became stronger during task performance (**Figure 4.1E**, right).

4.4.2 V1 neurons respond passively, whereas PPC responses are gated by task engagement.

To compare the overall response properties of V1 and PPC, we analyzed the responses of all neurons that had significant ($p < 0.01$, t-test) stimulus-period activity during the task. We focused on stimulus-period responses given that inactivation during this period disrupts behavioral performance (see Figure 3.3, Figure 3.5). A total of 1915 neurons (18% of all neurons, 18 fields, 9 mice) in V1 (**Figure 4.2A, B**) and 3524 neurons (26% of all neurons, 22 fields, 10 mice) in PPC (**Figure 4.2C, D**) were significantly responsive during the task. Two striking differences between V1 and PPC were immediately apparent when examining the trial-averaged responses. First, while V1 was about evenly split between target- and nontarget-preferring cells (61% target-preferring, **Figure 4.2E**), there was a strong bias in PPC towards the target stimulus (88% target-preferring). Secondly, while most task-responsive V1 neurons also responded during passive viewing (84% with significant passive response, $p < 0.01$, t-test, **Figure 4.2F**), the majority of PPC neurons had responses that were gated by task engagement (only 21% with significant passive response).

The dramatic effect of task engagement on PPC responses can also be observed by comparing the selectivity of responses during Passive and Engaged conditions (**Figure 4.2G**). We quantified selectivity using an ROC-based index that ranged from -1 to 1, with positive values indicating preference for the target stimulus (see Section 2.3.5 for details). V1 has a large proportion of significantly selective neurons in both Passive (83% of cells, $p < 0.05$, permutation test) and Engaged conditions (91% of cells). By contrast, PPC neurons are largely unselective during Passive conditions (31% selective, $p < 0.05$, permutation test), and yet become significantly selective to target trials during task engagement (91% of cells). We also quantified the change in responses for engagement versus passive viewing for each neuron using an engagement modulation index (**Figure 4.2H**), which also ranged from -1 to 1, with positive values

indicating increases with engagement. The mean modulation index was significantly above zero for target-preferring neurons in both V1 (0.23 ± 0.01; $p < 10^{-9}$, Wilcoxon signed rank test) and PPC (0.35 ± 0.01; $p < 10^{-9}$). Nontarget-preferring neurons in V1 showed weaker modulation (modulation index: 0.06 ± 0.01 for non-target; 0.23 ± 0.01 for target) than target-preferring neurons ($p < 10^{-9}$, Wilcoxon rank-sum test), indicating that the effect of engagement was stimulus-specific.

Comparing responses during engaged behavior and passive viewing therefore revealed different response properties in V1 and PPC. In V1, task performance did not merely increase overall responsiveness (as would be expected with arousal), but instead modulated firing rates to enhance the contrast between target and non-target stimuli. By contrast, PPC responses were strongly target-selective and gated by behavior. This dramatic task-dependent increase in activity strongly implicates a role for PPC beyond mere sensory processing. Furthermore, because this activity is selective to target trials in which the animal licks, a major component of PPC responses likely signals the impending action and does not simply reflect overall increases in arousal or attention during task engagement. A subset (~20%) of PPC neurons, however, do have significant passive responses, and could play a role in sensory processing.

4.4.3 Error trials reveal sensitivity of PPC to both stimulus and choice

One possible explanation for the strong behavioral gating and target selectivity of PPC responses is that these neurons reflect signals related to movement or action planning. Another possibility is that task engagement provides a stimulus-specific attentional signal. One way to arbitrate between these alternatives is to observe activity level on error trials. We thus examined PPC responses during error trials to see whether activity in PPC reflected the stimulus, the animal's eventual choice, or both.

We analyzed neuronal responses in V1 and PPC from behavioral sessions in which the mouse committed at least five misses or five false alarms. We focused on target-selective neurons,

as these represented the vast majority of responses in PPC. Because the mice had a bias towards licking, there were fewer sessions and neurons with sufficient Miss trials (V1, n = 202 cells from 3 sessions; PPC, 1157 cells from 8 sessions), compared to sessions with sufficient False Alarm trials (V1, n = 1053 cells from 16 sessions; PPC, n = 2996 cells from 21 sessions).

We compared neuronal responses on Miss trials to Hit trials as well as to trials with Passive presentation of target stimuli (**Figure 4.3A-B**). V1 responses were largely similar regardless of behavioral state (Miss versus Passive) or choice (Miss versus Hit). The majority of Hit-responsive V1 neurons were also significantly responsive during Miss trials (81%) and during Passive Target trials (71%). By contrast, many PPC neurons were strongly modulated by both behavioral state and choice. A majority of Hit-responsive PPC neurons (58%) did not show significant responses during either Miss or Passive trials. These neurons could be interpreted as encoding the animal's impending action, as they only responded on licking trials. Other PPC neurons (42%) did exhibit significant responses during Miss trials, although these cells were present only in a subset of imaged populations (4 of 6). These neurons were sensitive to behavioral state (Engaged versus Passive), and to the stimulus (Miss versus CR, not shown), but insensitive to the animal's choice (Hit versus Miss).

What happens on False Alarm (FA) trials, when mice lick inappropriately following a Nontarget stimulus? We compared neuronal responses on FA trials to both Hit trials and Correct Reject (CR) trials (**Figure 4.3C-D**). Hit-responsive V1 neurons varied in their responsiveness on Nontarget trials. A subset of V1 neurons (28%) responded significantly on CR trials, and a slightly larger proportion showed significant FA responses (44%). By contrast, PPC neurons showed much stronger responses on FA versus CR trials. A much larger proportion of neurons demonstrated significant FA responses (43%) than CR responses (5%), and this was true across imaged populations.

The above comparisons were made using responses measured during the stimulus period, and thus give a static snapshot of how responses differ on the various trial types. In a

complementary analysis, we also measured how the responses vary with trial type as a function of time. Such comparisons could then be used to evaluate the dynamics of stimulus and choice information encoded in the activity of V1 and PPC neurons (Goard et al., 2016).

We used an ROC-based analysis to compare single neuron responses on error trials with responses on correct trials at each time point of the trial (Britten et al., 1996; Hernandez et al., 2010). We assessed stimulus encoding by comparing Target trials (Hit, Miss) to Non-target trials (FA, CR), in which the same action was committed, but a different stimulus was presented. Conversely, we compared Lick trials (Hit, FA) to No-lick trials (Miss, CR) to assess choice encoding, where the same stimulus was presented, but a different action was chosen. We did this separately for sessions with sufficient Miss trials (**Figure 4.3E**) and for sessions with sufficient FA trials (**Figure 4.3H**). We quantified both the average selectivity across neurons (**Figure 4.3F,I**), and the proportion of neurons with significant selectivity (**Figure 4.3G,J**).

In V1, encoding of stimulus information was robust and appeared rapidly after stimulus presentation. A majority of neurons showed significant selectivity between Miss and CR trials (63%; **Figure 4.3G**, left), or between FA and Hit trials (83%; **Figure 4.3J**, left). Choice information was weak, with only a minority of neurons showing selectivity between Miss and Hit trials (13%) or between FA and CR trials (12%). For PPC, however, we observed robust encoding of both stimulus and choice (**Figure 4.3G, J**, right). Stimulus information was significant in a subset of PPC neurons for both Miss trials (37% selective) and FA trials (65%). Choice information appeared later in the trial compared to stimulus signals, but significant predictive choice activity was still observable during the stimulus period for a subset of neurons, whether looking at either Miss trials (21%) or FA trials (18%).

Although care must be taken to properly interpret differences in activity levels on error trials (see Discussion), PPC activity is neither purely sensory, as V1 responses appear to be, nor is it purely movement-related. Instead, PPC appears to be sensitive to both the stimulus and impending choice of the animal.

4.4.4 PPC reflects both stimulus contrast and behavioral state

Previous work has shown that in both primates (Gold and Shadlen, 2007; Shadlen and Newsome, 2001) and rodents (Hanks et al., 2015), neurons in PPC encode not only the impending motor action but also the sensory evidence for that decision. To test whether neurons in mouse PPC similarly reflected the decision process, we varied the amount of sensory evidence from trial-to-trial by manipulating stimulus contrast. We also compared responses during engaged and passive conditions to examine how sensory and motor signals may be combined in PPC responses.

A subset of the mice (n = 6) were trained to perform a variant of the discrimination task, in which the contrast of the grating stimulus varied randomly from trial-to-trial (**Figure 4.4A**). Mice performed well above chance, even at very low contrasts (d' at 2% contrast, 1.05 ± 0.25; $p < 0.05$, Wilcoxon signed-rank test), although performance degraded as contrast was lowered ($p < 0.05$, Kruskal-Wallis test) reflecting the decrease in the strength of sensory evidence (**Figure 4.4B**). We imaged from neurons in V1 (n = 8 sessions) and PPC (n = 11 sessions) and computed contrast response functions for both passive and engaged conditions. Because the large majority of neurons in PPC responded selectively to target stimuli (Figure 4.2C,F, 95% in this dataset), we focused our analysis on target-preferring neurons (V1, n = 250; PPC, n = 611). Neurons were included for further analysis if they demonstrated significant Hit responses at multiple contrasts that could be well fit with a hyperbolic ratio function (Albrecht and Hamilton, 1982) (see Section 4.3.8).

As in the single-contrast task, the population response in V1 was robust in both Engaged and Passive conditions (**Figure 4.5A**), with most neurons showing significant responses during both Engaged and Passive conditions (68% of cells, $p < 0.01$, t-test, **Figure 4.5B**). The majority of V1 neurons had Hit responses that were significantly modulated by contrast (73% of cells, $p < 0.05$, permutation test, **Figure 4.5C,E**, left), and a smaller proportion were significantly modulated

by engagement (41% of cells, $p < 0.05$, permutation test, **Figure 4.5D,E**, left). By contrast, PPC population activity was more robust in Engaged versus Passive conditions (**Figure 4.4C**). Only a subset of PPC neurons had have significant passive responses (24% of cells), though more neurons had Engaged responses that depended significantly on contrast (43% of cells, **Figure 4.5C,E**, right). The majority of PPC neurons were significantly enhanced by engagement (55% of cells, **Figure 4.5D,E**, right).

Are these sensory and motor signals encoded in separate neuron populations within PPC? Closer examination of the individual PPC contrast response functions revealed a great deal of heterogeneity (**Figure 4.4D**). We divided neurons into groups based on whether they exhibited significant modulation by contrast and/or engagement (**Figure 4.4D,E**; **Figure 4.5E**). A subset of PPC neurons (18% of cells) showed strong modulation by contrast, but very weak modulation by engagement (**Figure 4.4E**, left column). Such neurons therefore faithfully represented the sensory stimulus regardless of behavioral state. Conversely, a larger group of PPC neurons (29% of cells) were gated by task engagement but showed little to no modulation with contrast (**Figure 4.4B**, middle row). These neurons reflected the behavioral state and impending action of the animal irrespective of sensory drive. Lastly, a third subset of PPC neurons (25% of cells) were significantly modulated by both contrast and engagement (**Figure 4.4B**, bottom row). PPC therefore appears to contain both contrast-modulated "sensory" neurons as well as engagement-modulated "motor" neurons. This differs qualitatively from V1 (**Figure 4.5E**, left), where most neurons (73%) are modulated by contrast, and much fewer by engagement alone (11%).

4.4.5 PPC encodes both contrast-dependent sensory signals and contrast-independent choice signals

We also analyzed the error trials to see whether stimulus and choice signals could be separately extracted from responses in PPC, and whether these signals depended on contrast (**Figure 4.6**). We compared False Alarm trials with Hit and Correct Reject trials, using sessions with at least

five False Alarm trials at each contrast (7 of 11 sessions, n = 392 neurons). We did not make comparisons with Miss trials given the low number of trials. The population response on FA trials was weak, but distinguishable from the response on CR trials (**Figure 4.6A**), especially during the choice period (**Figure 4.6B**). Using an ROC-based approach, we again found that PPC encoded both stimulus and choice with differing time courses (**Figure 4.6E**). Interestingly, some PPC cells encoded the stimulus in a contrast-dependent manner, but also encoded the choice in a contrast-independent manner (**Figure 4.6C**). We quantified the contrast-dependence of the auROC index for each neuron by measuring its slope as a function of contrast (**Figure 4.6D**). A larger proportion of neurons exhibited significant contrast-dependence in stimulus encoding (**Figure 4.6F**, 15% of cells, $p < 0.05$, permutation test) compared to the proportion with significant contrast-dependence in choice encoding (4% of cells, $p < 0.05$). PPC may therefore simultaneously encode both contrast-dependent sensory signals and contrast-independent motor signals in the same population of neurons.

4.4.6 PPC reflects changes in stimulus-reward contingency

Thus far we have seen that PPC has a strong bias towards target stimuli during the task. Most neurons reflect the behavioral state, and a subset of neurons appear to encode the sensory stimulus in a contrast-dependent manner. Error analyses indicate that PPC neurons may encode both stimulus and choice signals, but errors can reflect other factors such as impulsivity or inattention. To more conclusively determine whether PPC neurons reflect sensory stimuli versus sensorimotor transformation, we manipulated the stimulus-reward structure of the task. By re-training mice on a reversed stimulus-reward contingency, and measuring activity from the same neurons before and after reversal, we could test whether neurons were more sensitive to stimulus identity or the learned task contingencies.

After imaging the responses of neurons in V1 and PPC in the original go/no-go task, we reversed the reward contingencies of the stimuli (**Figure 4.7A**). Licking in response to the original non-target stimulus (Stimulus B, blue) was now rewarded with water, whereas licking to Stimulus A (red) was punished with quinine. Three mice successfully learned the task after 7-11 days of training, although performance was slightly worse than before (**Figure 4.7B**; d-prime, original, 2.96 ± 0.50; reversed, 1.51 ± 0.22). We then measured responses from the same populations of neurons in V1 and in PPC under the reversed reward contingency, using a semi-automated procedure (Huber et al., 2012) to identify the same neurons across different imaging sessions (see Section 4.3.5 for details).

We analyzed the selectivity of individual neurons that had significant responses both before and after reversal (V1, n = 488 in 8 fields; PPC, n= 509 in 8 fields). Many neurons in V1 that were selective to a particular stimulus remained selective to the same stimulus after reversal, whether Stimulus A or Stimulus B (**Figure 4.7C**). By contrast, many PPC neurons exhibited a switch in selectivity when the contingency was reversed. These neurons that were selective to target stimulus A during the original task became selective to the new target stimulus B after reversal (**Figure 4.7D**).

We first examined the overall selectivity in neurons in V1 and PPC (**Figure 4.7E, F**). For V1, the representation of the two stimuli was relatively unchanged ($p = 0.13$, permutation test), with a slight bias towards the original target stimulus A. However, in PPC, selectivity was dramatically altered with reversal of reward contingency ($p < 0.001$, permutation test), with the majority of responsive neurons preferring the new target stimulus B after reversal.

We then plotted the selectivity of individual neurons before and after reversal against each other as a scatter plot (**Figure 4.7G,H**). Purely stimulus-selective neurons will remain close to the unity line and in the first (bottom-left) and third (top-right) quadrants, whereas neurons that are sensitive to the reward contingency will lie in either the second quadrant (top-left; for no-go selective cells) or the fourth quadrant (bottom-right; for go-selective cells). V1 neurons were

strongly stimulus-selective, with the majority (74% of cells, **Figure 4.8A**, top) of neurons lying within the first and third quadrants (**Figure 4.7G**). By contrast, the majority (67% of cells, **Figure 4.8A**, bottom) of PPC neurons were found in the fourth quadrant, indicating a preference for the rewarded target stimulus, regardless of its actual identity (**Figure 4.7H**). These results indicate that the majority of PPC neurons indeed reflect the learned task contingencies.

4.4.7 PPC cells with reversed selectivity exhibit stronger modulation by engagement and by error trials.

A small subset of PPC neurons (7%), however, did show stable stimulus selectivity before and after reversal. Given that a minority of neurons also exhibited stimulus-driven passive responses that were also insensitive to behavioral choice on error trials, we wondered whether these different properties (dynamic selectivity, engagement modulation, and error modulation) were related and could be used to separate functional cell types. In other words, if PPC neurons could indeed be separated into groups of "sensory" and "motor" neurons, then stimulus selectivity after reversal should be predictive of modulation strength with engagement and choice.

To test this hypothesis, we measured the strength of engagement modulation in both stimulus-selective (to either Stimulus A or B) neurons and go-selective neurons, as defined by their selectivity before and after reversal (**Figure 4.8A**). We found that go-selective PPC neurons had significantly higher modulation by engagement than stimulus-selective neurons (go-selective, 0.39 ± 0.01; stimulus-selective, 0.08 ± 0.04; $p < 10^{-9}$, Wilcoxon rank-sum test; **Figure 4.8B**, bottom). Indeed, many PPC neurons that maintained selectivity to Stimulus A after reversal also had strong passive responses (**Figure 4.8D**), and many neurons with reversed selectivity were strongly gated by engagement (**Figure 4.8E**). We also tested to see whether a PPC neuron's responses on error trials related to its selectivity after reversal. Go-selective PPC neurons had significantly higher error modulation (computed by comparing False Alarm trials to Correct Reject trials) compared to stimulus-selective neurons (go-selective, 0.27 ± 0.01; stimulus-selective, 0.02

± 0.02; $p < 10^{-9}$, Wilcoxon rank-sum test; **Figure 4.8C**). Interestingly, the small subset of go-selective neurons in V1 (6% of cells) also differed from stimulus-selective V1 neurons in their degree of engagement (**Figure 4.8B**, top) and error (**Figure 4.8C**, top) modulation, suggesting that these different properties are strongly related and can be used to distinguish functional cell types across different areas.

These findings demonstrate distinct subsets of neurons within PPC. One subset of neurons faithfully reflects the sensory input, both in passive conditions and after learning a new reward contingency. The larger proportion of PPC neurons, however, switch selectivity after re-training, and are strongly modulated by behavioral state. The flexibility and heterogeneity of PPC responses suggests a role for PPC in the mapping of sensory inputs onto appropriate motor actions.

4.5 Discussion

We developed a head-fixed visual decision task for mice with separate stimulus and response epochs, and used population imaging to investigate the role of PPC in perceptual decisions. Our key findings are that PPC encodes both sensory and motor signals across a heterogeneous pool of neurons, and that its activity can change drastically depending on task performance and demands. Together these results suggest that mouse PPC is responsible for neither pure sensory processing, nor for the control of motor output, but rather is important for the decision process itself – the process of mapping sensation to action.

4.5.1 Mouse PPC is more than an extrastriate visual area

The small size of the mouse brain has made it difficult to identify precise borders between different functional areas. PPC in the rodent, as classically defined by its thalamic inputs (Reep et al., 1994), is located in the region between the more posterior V1 and the more anterior

somatosensory cortex. However, this location is essentially where both anatomical (Wang and Burkhalter, 2007; Wang et al., 2012) and functional (Garrett et al., 2014; Marshel et al., 2011) mapping studies have identified the retinotopically-organized areas RL, A, and AM. The degree to which PPC overlaps with these secondary visual areas is a matter of debate. Some have argued that rodent PPC may have more in common with primate extrastriate cortex, given that inactivation specifically disrupts sensitivity on visual decisions (Licata et al., 2016). Indeed, the stereotaxic coordinates used by us (Goard et al., 2016) and others (Harvey et al., 2012; Morcos and Harvey, 2016) to target mouse PPC most directly overlap with area AM, which exhibits directionally-tuned responses even under anesthesia (Marshel et al., 2011). But is PPC simply a sensory visual area?

We have presented multiple pieces of evidence that point to a role for PPC beyond mere sensory processing. First, activity in PPC is strongly dependent on behavioral state, with only a minority of task-responsive PPC neurons exhibiting significant responses during passive viewing of stimuli (Figure 4.2F). Secondly, the selectivity of PPC neurons is strongly biased toward target stimuli (Figure 4.2E). This bias is likely due to the asymmetry of the Go/No-go paradigm, and may reflect a learned association of stimulus and reward (Fitzgerald et al., 2013). Third, PPC responses are modulated during error trials (Figure 4.3, Figure 4.6). Information about the eventual choice of the animal can be decoded from the activity of PPC, as previously shown in both mice (Goard et al., 2016; Harvey et al., 2012) and rats (Raposo et al., 2014). Finally, and most conclusively, the biased selectivity of most PPC neurons towards target stimuli is dramatically reversed when the animal is retrained on a different reward contingency (Figure 4.7). Together these results demonstrate that the stimulus-period responses in PPC are task-dependent and not purely sensory.

4.5.2 PPC activity does not merely reflect movement

One may argue that the activity patterns we observed in PPC can be most parsimoniously explained as movement or action-planning related signals. After all, movement-related activity would be present only during engagement, it would exhibit choice selectivity, and it would change with reward contingency. It is unlikely that PPC is directly involved in executing motor plans, as we have previously shown that optogenetic inactivation of PPC during the response period, or even during the delay between stimulus and response, has no effect on behavior (Goard et al., 2016). However, the possibility remains that the PPC responses recorded in our task reflect planning- or movement-related signals that originate elsewhere. Although some PPC neurons (29%) do show activity that appears motor-related, due to their contrast-independent nature (Figure 4.4E), we provided evidence that PPC is not a purely motor area either.

First, passive visual stimulation does induce a response in some PPC neurons (20%), as previously shown in parietal area AM of anesthetized mice (Marshel et al., 2011). This subset of neurons tends to stably reflect the stimulus even with changes in reward contingency (Figure 4.8). Second, we find a subset of PPC neurons that reflect the target stimulus on Miss trials when the animal fails to lick, even though such neurons are inactive during passive viewing (Figure 4.3). Finally, the responses of many PPC neurons (43%) shows modulation with stimulus contrast, even for the same decision and motor output (Figure 4.4E). This is reminiscent of previous primate (Shadlen and Newsome, 2001) and rodent (Hanks et al., 2015) PPC studies, where it has been shown that responses vary with the strength of incoming sensory evidence.

4.5.3 Effects of attention and behavioral state

The posterior parietal cortex has been implicated not only in decision-making, but also in spatial attention and visual salience (Colby and Goldberg, 1999). To what degree could our results be explained by modulation of attention, arousal, or behavioral state? PPC responses were strongly

target-selective, but the reversal experiments demonstrate that this selectivity was not stimulus-specific but instead reflected a preference for Go trials. These results argue against the possibility of a stable sensory representation in PPC that is enhanced by attentional engagement. Nonetheless, we cannot fully distinguish whether the Go-selective signals in PPC are primarily related to the animal's choice (whether decision formation or motor planning) or instead to trial-specific changes in attention or arousal. For example, differences of response amplitude on Go versus No-go trials could be mediated by neuromodulatory input in PPC that reflects variations in arousal or attentional effort (Sarter et al., 2005). Similarly, the variation of PPC responses with contrast could also be explained by changes in attention, rather than by different levels of evidence for a decision. To better isolate the decision process from attention, future work should utilize forced-choice designs which are more immune to differences in motivation and arousal (Carandini and Churchland, 2013), and which are also becoming more amenable for use in rodents (Brunton et al., 2013; Guo et al., 2014a; Raposo et al., 2014).

4.5.4 Heterogeneity of PPC responses

In every task condition reported here, PPC responses were heterogeneous. A subset of PPC neurons have significant passive visual responses and are modulated by contrast in both engaged and passive conditions. These neurons maintain their stimulus selectivity even after reversal of reward contingency. These "sensory" neurons are spatially intermingled with the larger proportion of "motor" neurons that have task-gated responses, weak contrast modulation, and altered selectivity after contingency reversal.

Heterogeneous response properties have been previously reported in both primate (Bennur and Gold, 2011; Meister et al., 2013) and rat (Raposo et al., 2014) PPC during decision tasks. Our work adds to this literature, and additionally provides evidence that such

heterogeneous responses exist in a spatially intermingled fashion within PPC of the mouse. Do these response types form bona-fide cell classes or does PPC represent a category-free population, as others have proposed (Raposo et al., 2014)? Although we do find that various properties (such as modulation by engagement, contrast, and reversal) are correlated with one another, more work needs to be done to determine whether functional neuronal subgroups are truly separable. Future work can take advantage of the tools available for the mouse system to determine whether these functional properties correlate with molecular markers of cell identity (Kvitsiani et al., 2013; Pinto and Dan, 2015) or with projection target (Chen et al., 2013a; Li et al., 2015). Additionally, while we did not observe clear spatial clustering of response types in our experiments, anatomical experiments have demonstrated clear differences in connectivity for the medial versus lateral subdivisions of PPC (Wilber et al., 2014). Finer delineation of areal boundaries, perhaps using retinotopic mapping (Garrett et al., 2014) or noise correlation analyses (Kiani et al., 2015), could help distinguish whether functional properties vary systematically with anatomical location within the parietal cortex.

4.5.5 Conclusion and outlook

Our results demonstrate that mouse PPC encodes sensory, decision, and motor variables, suggesting a role in sensorimotor transformation during perceptual decisions. Our reversal experiments reveal a strong experience-dependent plasticity of representations in PPC. Although previous studies have similarly shown that representations in PPC can be altered with re-training (Freedman and Assad, 2006), our work unequivocally shows such changes in individual PPC neurons. Taking advantage of the stability and consistency of two-photon imaging, we demonstrate that not only PPC as a whole, but also many individual neurons show a remarkable switch in selectivity after reversal of the task reward structure.

The mouse posterior parietal cortex encodes behaviorally-relevant variables in a highly task-dependent manner, in analogy to prior work in primates. Our understanding of how decisions are computed and sensorimotor transformations are made will be greatly aided by future circuit-level analyses of PPC function in this powerful model system (Carandini and Churchland, 2013).

4.6 Contributions and Acknowledgements

G.N.P. and M.J.G. designed experiments with input from M.S.; M.J.G. performed the surgeries; M.J.G. and G.N.P. performed the imaging experiments; M.J.G., G.N.P., J.W., and B.C. performed the behavioral training; G.N.P. developed the behavior software; G.N.P. analyzed the data with comments from M.J.G. and M.S.; G.N.P., M.J.G., and M.S. wrote the manuscript.

We thank J. Sharma, A. Boesch, V. Li, C. Le, J. Krizan, V. Breton-Provencher, and T. Emery for technical assistance; E. K. Miller and R. Desimone for discussions of the behavioral task; R. Huda, M. Hu, H. Sugihara and C. D. Harvey for helpful discussions and/or comments on the manuscript; L. L. Looger, J. Akerboom, D. S. Kim, and the Genetically-Encoded Neuronal Indicator and Effector (GENIE) Project at Janelia Farm Research Campus Howard Hughes Medical Institute for generating and characterizing GCaMP6 variants. This work was supported by the NIH (M.J.G., F32-EY023523 and K99-MH104259; M.S., R01-EY007023 and U01-NS090473); the NSF (G.N.P., Graduate Research Fellowship; M.S., EF1451125); the Simons Center for the Social Brain (M.S.), and the Picower Institute Innovation Fund (M.J.G.; M.S.).

4.7 References

Albrecht, D.G., and Hamilton, D.B. (1982). Striate cortex of monkey and cat: contrast response function. J Neurophysiol *48*, 217-237.

Andermann, M.L., Kerlin, A.M., and Reid, R.C. (2010). Chronic cellular imaging of mouse visual cortex during operant behavior and passive viewing. Frontiers in cellular neuroscience *4*, 3.

Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. J Neurosci *31*, 913-921.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Visual neuroscience *13*, 87-100.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95-98.

Carandini, M., and Churchland, A.K. (2013). Probing perceptual decisions in rodents. Nat Neurosci *16*, 824-831.

Carandini, M., Heeger, D.J., and Movshon, J.A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. J Neurosci *17*, 8621-8644.

Chen, J.L., Carta, S., Soldado-Magraner, J., Schneider, B.L., and Helmchen, F. (2013a). Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. Nature *499*, 336-340.

Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V.*, et al.* (2013b). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295-300.

Colby, C.L., and Goldberg, M.E. (1999). Space and attention in parietal cortex. Annu Rev Neurosci *22*, 319-349.

Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. Elife *4*.

Fitzgerald, J.K., Freedman, D.J., Fanini, A., Bennur, S., Gold, J.I., and Assad, J.A. (2013). Biased associative representations in parietal cortex. Neuron *77*, 180-191.

Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. Nature *443*, 85-88.

Garrett, M.E., Nauhaus, I., Marshel, J.H., and Callaway, E.M. (2014). Topography and areal organization of mouse visual cortex. J Neurosci *34*, 12587-12600.

Goard, M.J., Pho, G.N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife *5*, 558.523.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annual review of neuroscience *30*, 535-574.

Green, D.M., and Swets, J.A. (1966). Signal detection theory and psychophysics, Vol 1974 (Wiley New York).

Guo, Z.V., Hires, S.A., Li, N., O'Connor, D.H., Komiyama, T., Ophir, E., Huber, D., Bonardi, C., Morandell, K., Gutnisky, D.*, et al.* (2014a). Procedures for behavioral experiments in head-fixed mice. PloS one *9*, e88678.

Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014b). Flow of cortical activity underlying a tactile decision in mice. Neuron *81*, 179-194.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220-223.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62-68.

Hernandez, A., Nacher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vazquez, Y., Camarillo, L., and Romo, R. (2010). Decoding a perceptual decision process across cortex. Neuron *66*, 300-314.

Huber, D., Gutnisky, D.A., Peron, S., O'Connor, D.H., Wiegert, J.S., Tian, L., Oertner, T.G., Looger, L.L., and Svoboda, K. (2012). Multiple dynamic representations in the motor cortex during sensorimotor learning. Nature *484*, 473-478.

Kiani, R., Cueva, C.J., Reppas, J.B., Peixoto, D., Ryu, S.I., and Newsome, W.T. (2015). Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. Neuron *85*, 1359-1373.

Kleiner, M., Brainard, D., and Pelli, D. (2007). What's new in Psychtoolbox-3? Perception *36*, 14-14.

Kvitsiani, D., Ranade, S., Hangya, B., Taniguchi, H., Huang, J.Z., and Kepecs, A. (2013). Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. Nature *498*, 363-366.

Li, N., Chen, T.W., Guo, Z.V., Gerfen, C.R., and Svoboda, K. (2015). A motor cortex circuit for motor planning and movement. Nature *519*, 51-56.

Licata, A.M., Kaufman, M.T., Raposo, D., Ryan, M.B., Sheppard, J.P., and Churchland, A.K. (2016). Posterior parietal cortex guides visual decisions in rats. bioRxiv.

Marshel, J.H., Garrett, M.E., Nauhaus, I., and Callaway, E.M. (2011). Functional specialization of seven mouse visual cortical areas. Neuron *72*, 1040-1054.

Meister, M.L., Hennig, J.A., and Huk, A.C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. J Neurosci *33*, 2254-2267.

Morcos, A.S., and Harvey, C.D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. Nat Neurosci *19*, 1672-1681.

Pinto, L., and Dan, Y. (2015). Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed Behavior. Neuron *87*, 437-450.

Pinto, L., Goard, M.J., Estandian, D., Xu, M., Kwan, A.C., Lee, S.H., Harrison, T.C., Feng, G., and Dan, Y. (2013). Fast modulation of visual perception by basal forebrain cholinergic neurons. Nat Neurosci *16*, 1857-1863.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat Neurosci *17*, 1784-1792.

Reep, R.L., Chandler, H.C., King, V., and Corwin, J.V. (1994). Rat posterior parietal cortex: topography of corticocortical and thalamic connections. Experimental brain research *100*, 67-84.

Sarter, M., Hasselmo, M.E., Bruno, J.P., and Givens, B. (2005). Unraveling the attentional functions of cortical cholinergic inputs: interactions between signal-driven and cognitive modulation of signal detection. Brain Res Brain Res Rev *48*, 98-111.

Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. Journal of neurophysiology *86*, 1916-1936.

Wang, Q., and Burkhalter, A. (2007). Area map of mouse visual cortex. The Journal of comparative neurology *502*, 339-357.

Wang, Q., Sporns, O., and Burkhalter, A. (2012). Network analysis of corticocortical connections reveals ventral and dorsal processing streams in mouse visual cortex. J Neurosci *32*, 4386-4399.

Wilber, A.A., Clark, B.J., Demecha, A.J., Mesina, L., Vos, J.M., and McNaughton, B.L. (2014). Cortical connectivity maps reveal anatomically distinct areas in the parietal cortex of the rat. Front Neural Circuits *8*, 146.

4.8 Figures

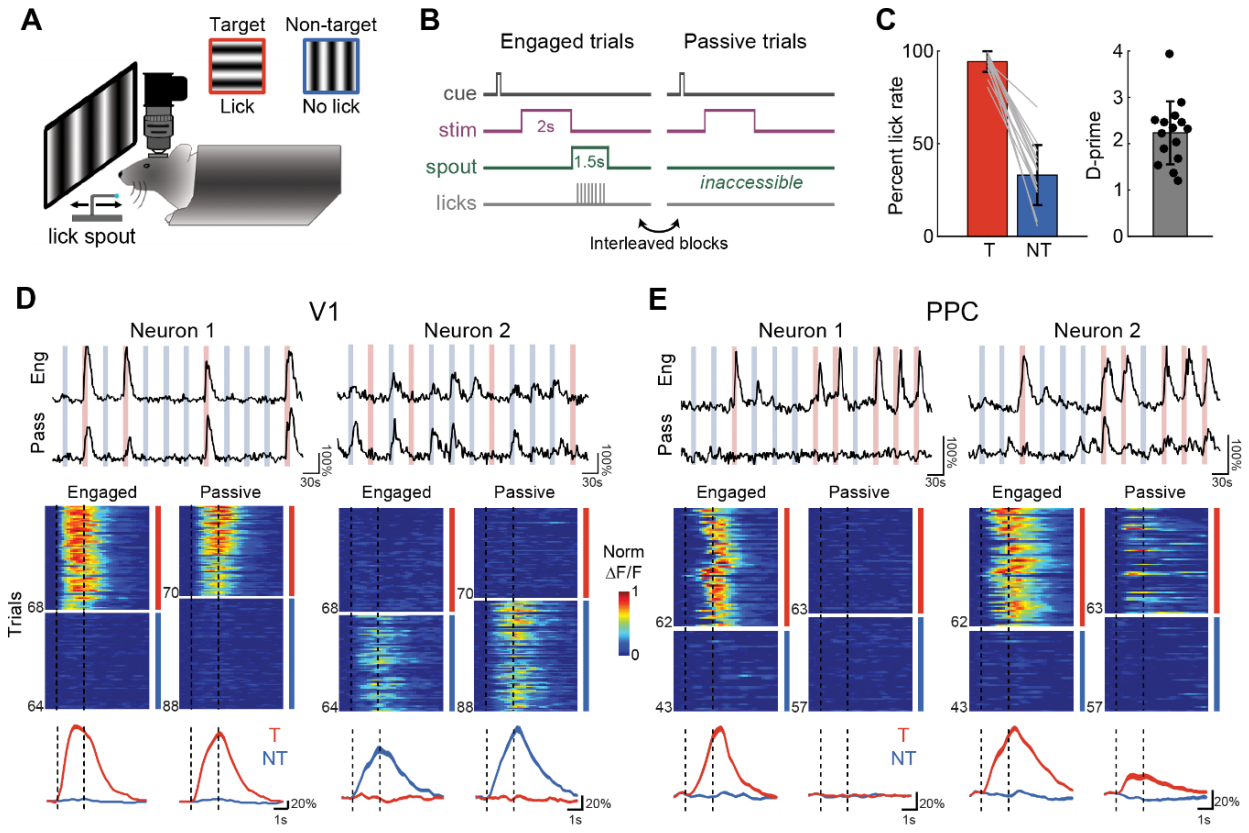

Figure 4.1. Imaging calcium responses in V1 and PPC during engaged task performance and passive viewing.

(A) Head-fixed mice were trained to perform a go, no-go lick-based visual discrimination task. A retractable lick spout was used to restrict lick responses to a specific epoch of the task. Licks following a target stimulus (red, horizontal drifting upwards, Stimulus A) were rewarded with water, while licks to non-target stimulus (blue, vertical drifting rightwards, Stimulus B) were punished with quinine. (**B**) Trial structure for Engaged and Passive conditions. After a brief auditory preparatory cue, a drifting grating was presented for 2 s. During Engaged trials, the retractable lick spout was presented immediately after stimulus offset for a minimum of 1.5 s. During Passive trials, the spout was inaccessible. Engaged and Passive trials were presented in blocks which were usually interleaved. (**C**) Rate of licking on target (T, red) and non-target (NT, blue) for each mouse using in imaging experiments (n=15). Behavioral performance was quantified as d-prime (mean across mice: 2.22). Error bars in this and all subsequent figures depict mean ± SEM. (**D**) Stimulus-evoked response of two V1 neurons, one target-selective (left column), and one non-target selective (right). Top, raw calcium response to multiple presentations of target (red) and non-target (blue) stimuli during both Engaged and Passive conditions. Middle, heatmap of trial-to-trial responses to target and non-target stimuli, presented in alternating blocks of Engaged (left) and Passive (right) trials, normalized to max response. Vertical dashed lines demarcate duration of stimulus. Bottom, overlay of trial-averaged responses for each stimulus during Engaged (left) and Passive (right) conditions. Line thickness indicates mean ± SEM. (**E**) Same as (D) but for two PPC neurons.

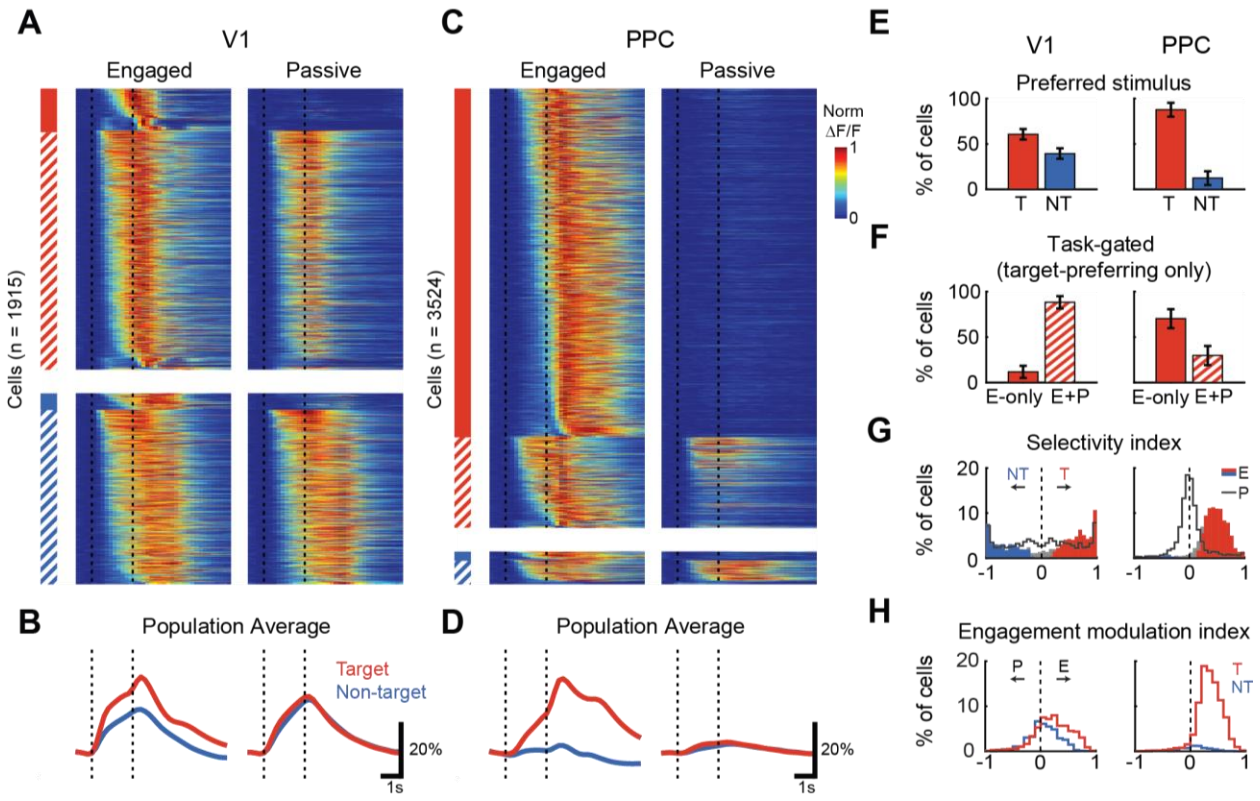

Figure 4.2. V1 neurons respond passively, whereas PPC responses are gated by task engagement.

(**A**) Trial-averaged responses of all task-responsive V1 neurons (n = 1915). Heatmap of all trial-averaged responses (preferred stimulus only) in both Engaged (left) and Passive (right) conditions, normalized by peak response. Neurons are separated into target-preferring (red) and non-target preferring (blue), and then into passive-responding (solid) and task-gated (hatched). Vertical dashed lines demarcate duration of stimulus. (**B**) Bottom, average response across V1 neurons to the target (red) and non-target (blue) in Engaged (left) and Passive (right) conditions. Line thickness indicates mean ± SEM. (**C-D**) Same as (A-B), but for PPC neurons (n = 3524). (**E**) Percentage of neurons in V1 (left) and PPC (right) that prefer target (T) or non-target (NT) stimulus. Error bars represent bootstrapped SEM across sessions. (**F**) Percentage of target-preferring neurons in V1 (left) and PPC (right) that are task-gated, i.e. that respond only during engagement (E-only), or that respond during both engaged and passive conditions (E+P). (**G**) Histogram of stimulus selectivity index for V1 (left) and PPC (right) during Engaged (filled bars) and Passive conditions (gray line). Positive selectivity indicates preference for Target. Colored bars indicate neurons with significant individual selectivity during Engaged trials. (**H**) Histogram of engagement modulation index for V1 (left) and PPC (right) for target- (red) and nontarget-preferring (blue) neurons.

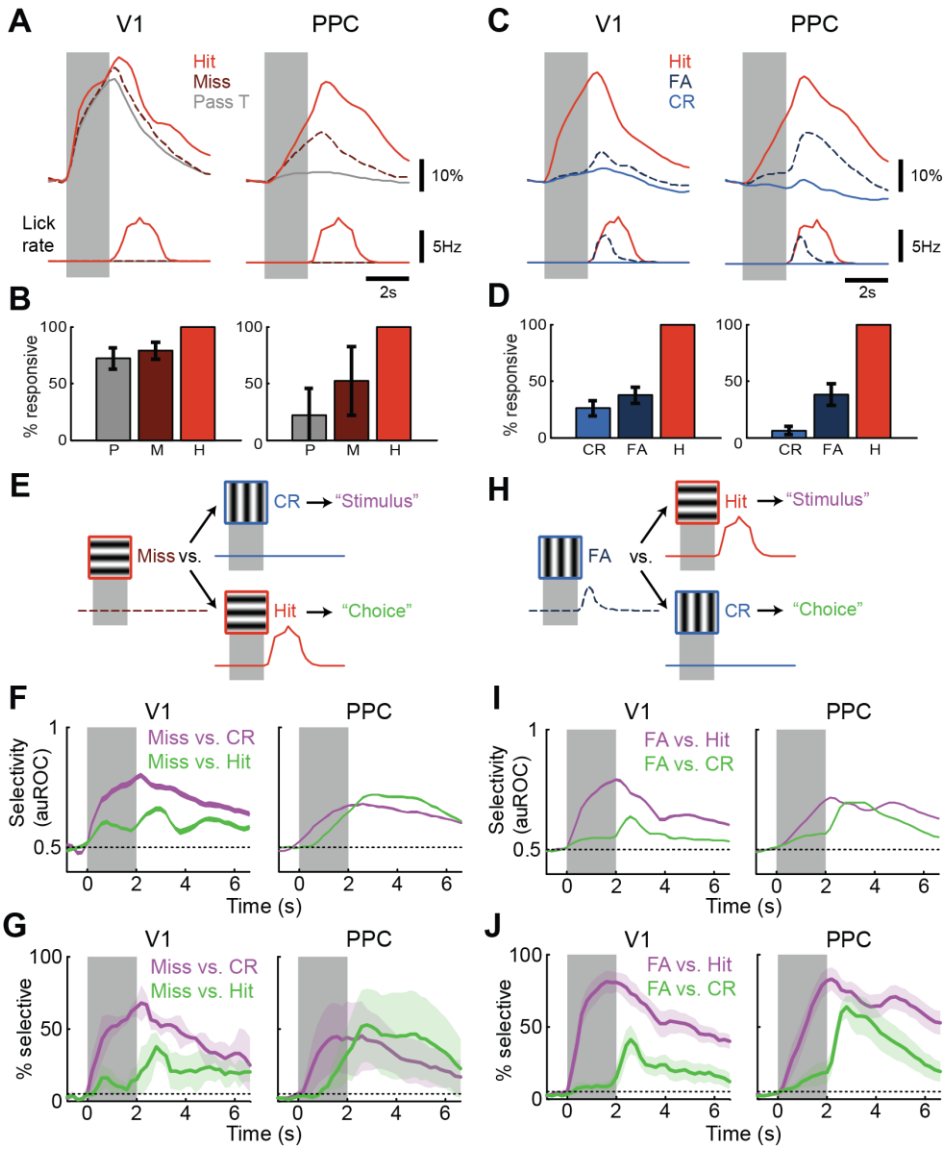

Figure 4.3. Error trials reveal sensitivity of PPC to both stimulus and choice.

(**A**) Population trial-averaged responses (top) during Miss trials, for recordings with at least five Miss trials, in V1 (left; 202 neurons across 3 sessions) and in PPC (right; 1157 neurons across 8 sessions). Average lick rate (bottom). (**B**) Percentage of neurons with significant responses on Passive Target (P), Miss (M), and Hit (H) trials for V1 (left) and PPC (right). Error bars represent bootstrapped SEM across sessions. (**C**) Population trial-averaged responses (top) during False Alarm (FA) trials, for recordings with at least five FA trials, in V1 (left; 1053 neurons across 16 sessions) and in PPC (right; 2996 neurons across 21 sessions). Average lick rate (bottom). (**D**) Percentage of neurons with significant responses on CR, FA and Hit trials for V1 (left) and PPC (right). (**E**) Miss trials are compared with CR trials to assess "Stimulus" encoding (purple), and are compared with Hit trials to assess "Choice" encoding (green). (**F-G**) Average selectivity (**F**) and percentage of selective neurons (**G**), as measured using a frame-by-frame ROC analysis to compare trial types, quantified as area under the ROC curve (auROC). Line thickness indicates mean ± SEM. Light gray shaded regions demarcate duration of stimulus. Dotted line indicates the expected performance/percentage by chance. (**H**) FA trials are compared with Hit trials to assess "Stimulus" encoding (purple), and are compared with CR trials to assess "Choice" encoding (green). (**I-J**) Same as (**F-G**) but for comparisons with FA trials.

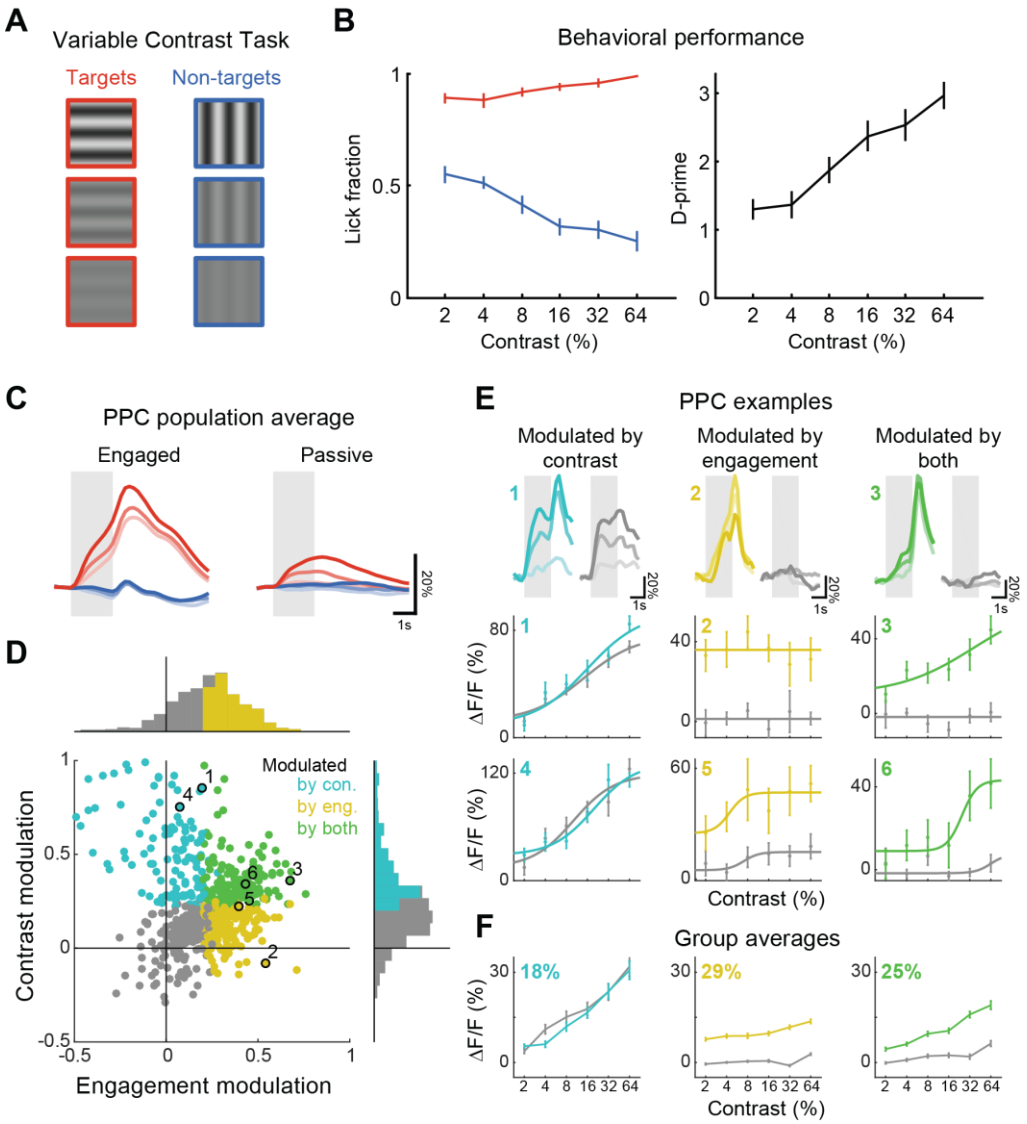

Figure 4.4. PPC reflects both stimulus contrast and behavioral state.

(**A**) Mice discriminated between orthogonally oriented Target and Non-target stimuli with contrast varying from 2 to 64%. (**B**) Behavioral performance (d-prime) on the variable contrast discrimination task, averaged across 19 sessions from 6 mice. Error bars indicate SEM. (**C**) PPC population trial-averaged response across contrasts during correct Engaged (left) and Passive (right) target (red) and nontarget (blue) trials. Responses are averaged across low (2 or 4%, light shade), medium (8 or 16%, medium shade), and high (32 or 64%, dark shade) contrast. Light gray shaded regions demarcate duration of stimulus. (**D**) Scatter plot and histograms of contrast modulation versus engagement modulation for all target-selective PPC neurons (n = 611) imaged during the variable contrast task. Colored bars in histograms (see also Figure 4.5C,D) indicate neurons with significantly positive modulation by contrast (cyan) or engagement (yellow). Colored dots on scatter plot demarcate neurons with significantly positive modulation by contrast alone (cyan), engagement alone (yellow), or both contrast and engagement (green). Individual examples in (**E**) are marked with the corresponding number. (**E**) Trial-averaged responses (top row) and contrast-response functions (middle and bottom rows) of example PPC neurons that were significantly modulated by contrast (left column), engagement (middle column), or both (right column). Modulation index values for each example can be found by referring to (**D**). (**F**) Group-averaged contrast-response functions. Percentages indicates proportion of PPC neurons within each group.

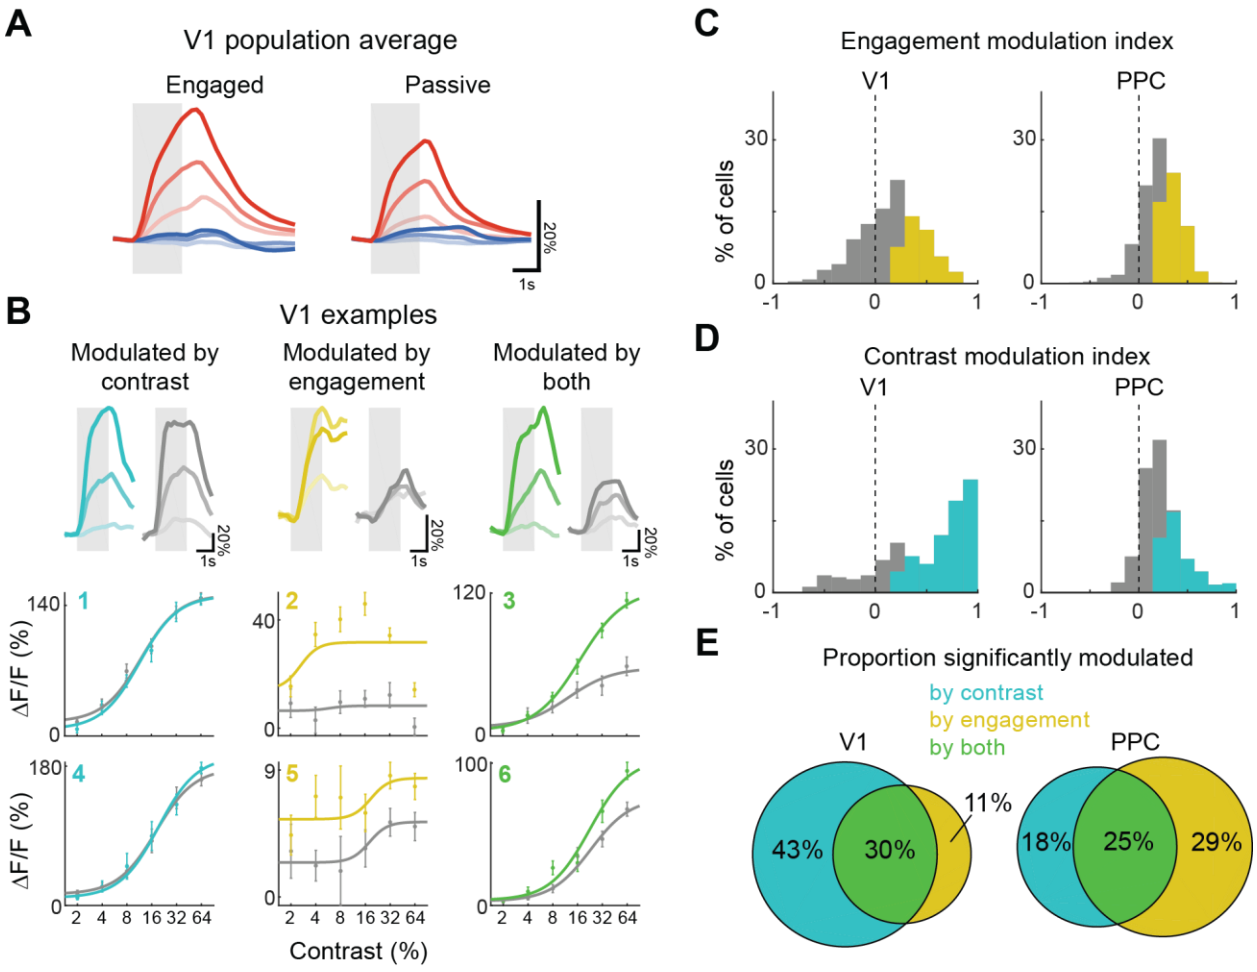

Figure 4.5. Comparison with V1 contrast responses.

(**A**) V1 population (n = 250 target-selective cells) trial-averaged response across contrasts during Engaged (left) and Passive (right) target (red) and nontarget (blue) trials. Responses are averaged across low (2 or 4%, light shade), medium (8 or 16%, medium shade), and high (32 or 64%, dark shade) contrast. Only correct trials are shown for Engaged condition. Light gray shaded regions demarcate duration of stimulus. (**B**) Trial-averaged responses (top row) and contrast-response functions (middle and bottom rows) of example V1 neurons that were significantly modulated by contrast (left column), engagement (middle column), or both (right column). (**C**) Histograms comparing engagement modulation in V1 and PPC, computed by comparing responses on Engaged versus Passive high contrast trials. Colored bars indicate neurons with significant modulation. (**D**) Histograms comparing contrast modulation in V1 and PPC, computed by comparing responses on high versus low contrast Engaged trials. Colored bars indicate neurons with significant modulation. (**E**) Venn diagrams indicating proportions of target-selective neurons in V1 and PPC with significantly positive modulation by contrast alone (cyan), engagement alone (yellow), or both contrast and engagement (green).

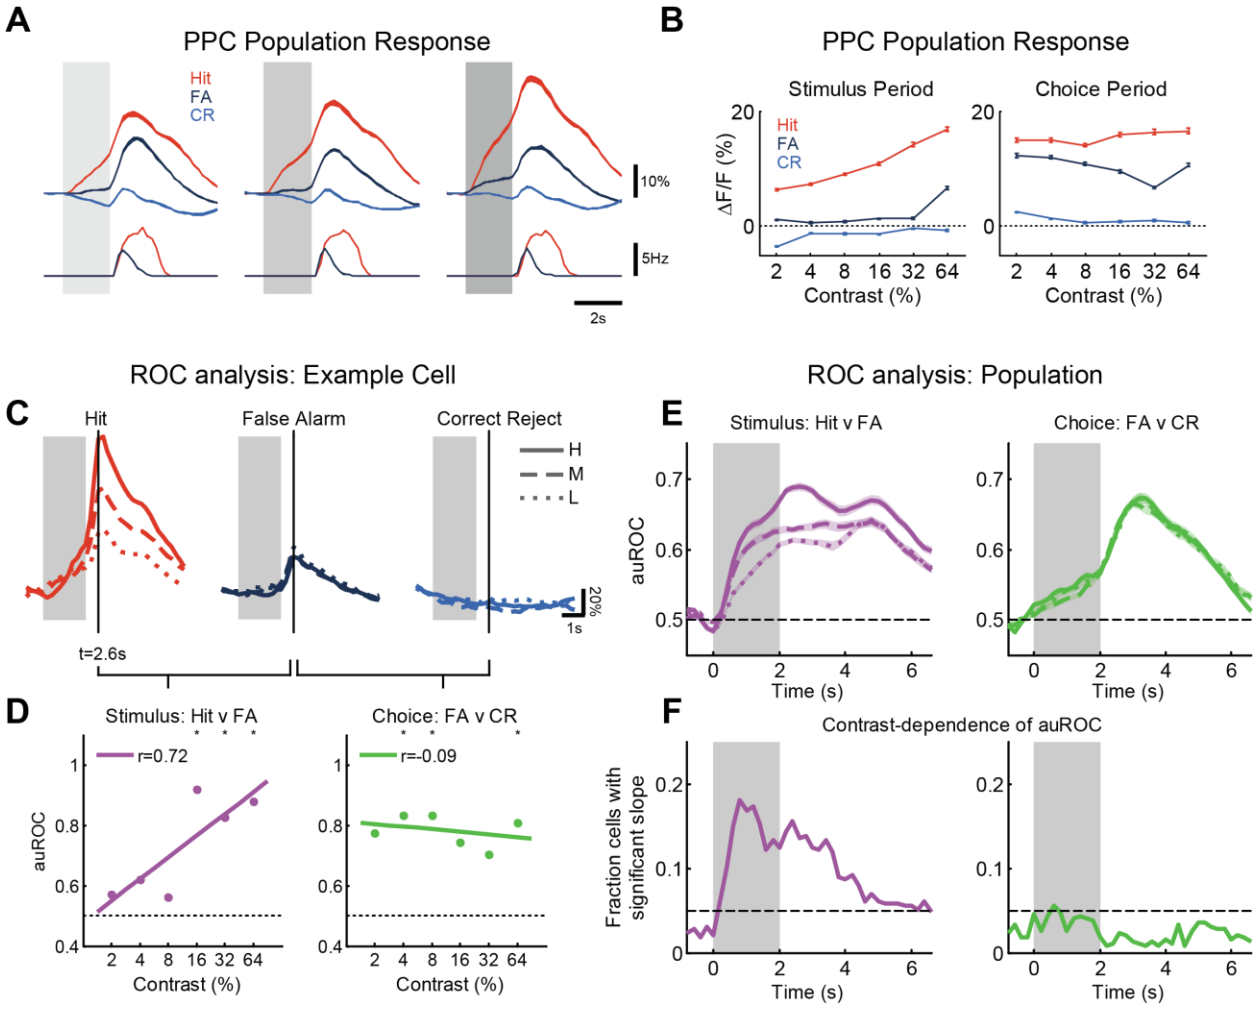

Figure 4.6. PPC encodes both contrast-dependent sensory signals and contrast-independent choice signals.

(**A**) Population average of PPC responses across contrasts during Hit, False Alarm (FA), and Correct Reject (CR) trials (top). Responses are averaged across low (2 or 4%, left), medium (8 or 16%, middle), and high (32 or 64%, right) contrast. Average lick rate (bottom). Light gray shaded regions demarcate duration of stimulus. (**B**) Time-averaged population response as a function of contrast during the Stimulus period (left, 0 to 2 s) or during the Choice period (right, 2 to 3.5 s) for Hit, FA, and CR trials. Choice period responses were measured relative to preceding Stimulus period response. (**C**) Response of an example PPC neuron during Hit trials (left), False Alarm trials (middle), and Correct Reject trials (right), across different contrasts, from low (L, dotted) to medium (M, dashed), to high (H, solid). (**D**) False Alarm (FA) trials were compared with Hit trials to assess Stimulus encoding (left), and with CR trials to assess Choice encoding (right) at a single time-point but across multiple contrasts. Best-fit line across contrasts is plotted along with its slope. At this time-point (2.6 s after stimulus onset), this neuron has significant (*, $p<0.05$) stimulus and choice encoding at multiple contrasts. Stimulus encoding is contrast-dependent, as seen with the significant positive slope, whereas choice encoding is not. (**E**) Average stimulus encoding (left) and choice encoding (right) across all PPC neurons as a function of time and of contrast, from low (L, dotted) to medium (M, dashed), to high (H, solid). Shading indicates SEM. (**F**) Fraction of PPC neurons with significant contrast-dependence in stimulus encoding (left) or choice encoding (right), as measured by slope of auROC with respect to contrast. Dashed line indicates fraction expected by chance.

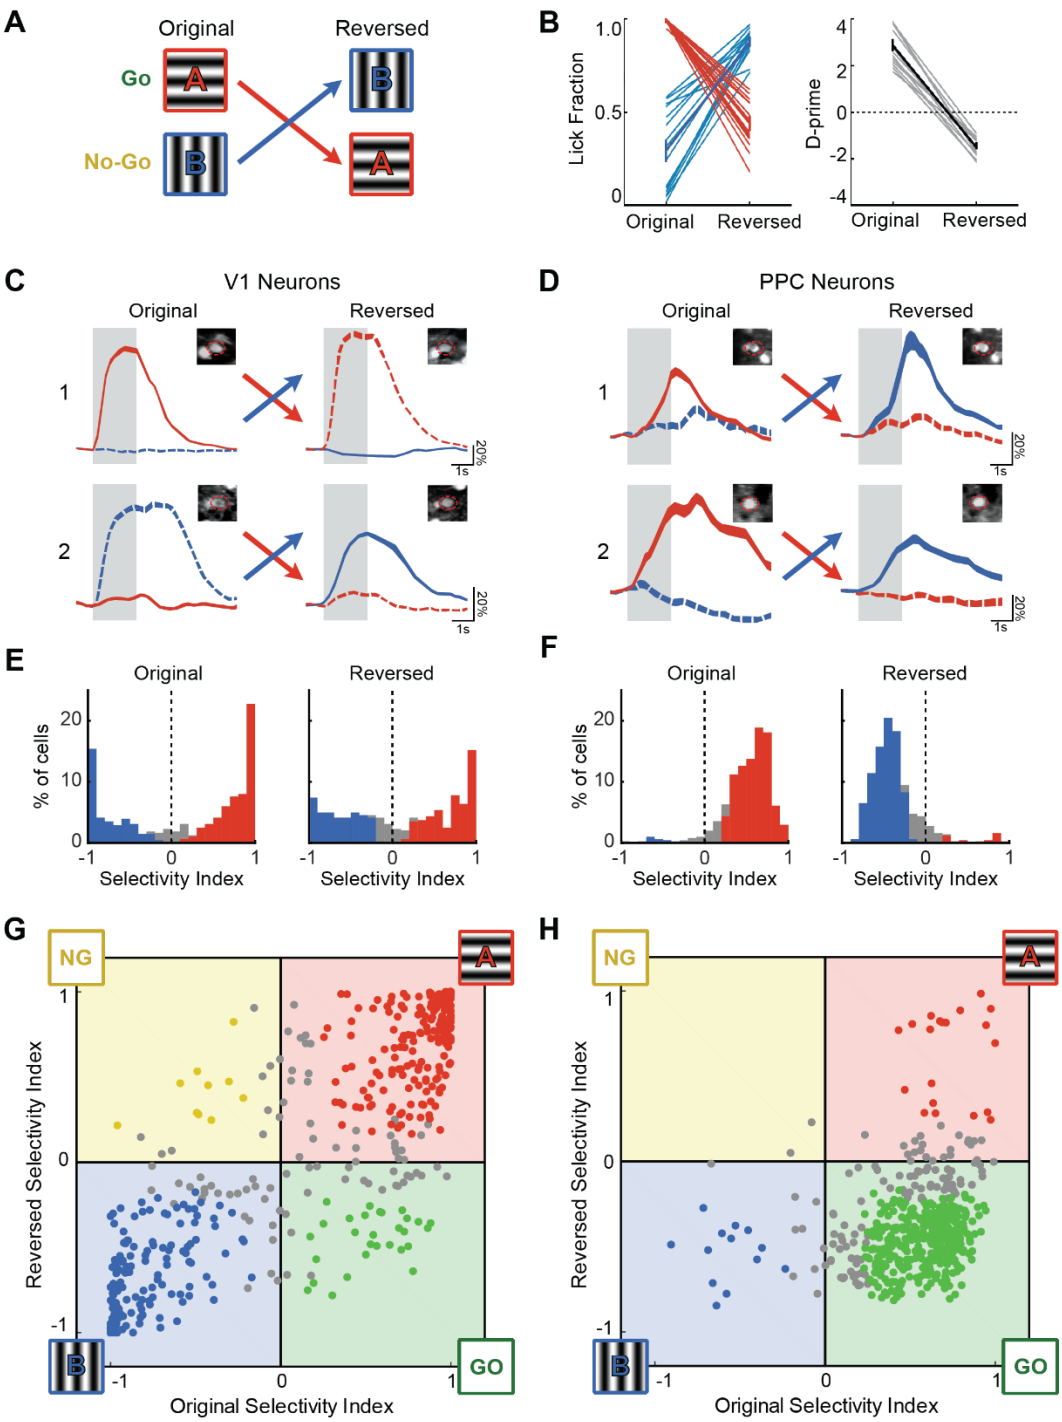

Figure 4.7. PPC reflects changes in stimulus-reward contingency.

(**A**) Mice were trained on a go/no-go discrimination task with a reversed reward contingency. (**B**) Behavioral performance before and after reversal. Left, response rate for Stimulus A (red) and Stimulus B (blue). Right, d-prime assuming Stimulus A as target. Performance after reversal was significantly different from zero ($p < 0.001$, Wilcoxon signed-rank test). (**C**) Neural response of two example V1 neurons before and after reversal of reward contingency. Colors indicate stimulus (red, A; blue, B) and line styles indicate whether the stimulus is target (solid) or non-target (dashed). Top row, response of a neuron selective to Stimulus A

both before (left) and after (right) reversal. Bottom row, response of a neuron selective to Stimulus B. Insets indicate the average projection image before and after reversal, and ROI (dashed red line) that was used to extract the neural response. (**D**) Same as (C) but for two PPC neurons. These neurons show a switch in selectivity with reversal of reward contingency. (**E**) Histogram of selectivity index before and after reversal for V1 neurons with significant responses both before and after reversal (n=488). Positive values indicate selectivity to the original Target, Stimulus A (red). Colored bars indicated neurons with significant selectivity. (**F**) Same as (**E**) but for PPC neurons (n=509). (**G**) Scatter plot comparing selectivity index in V1 before and after reversal. Colored points indicate neurons with significant selectivity both before and after reversal. Neurons in the first and third quadrants are stimulus-selective, as they prefer Stimulus A (red) or Stimulus B (blue) both before and after reversal. Neurons in the second and third quadrants are choice-selective, as they prefer either the Go stimulus (green) or the No-Go stimulus (yellow) in both conditions. (**H**) Same as (G), but for PPC.

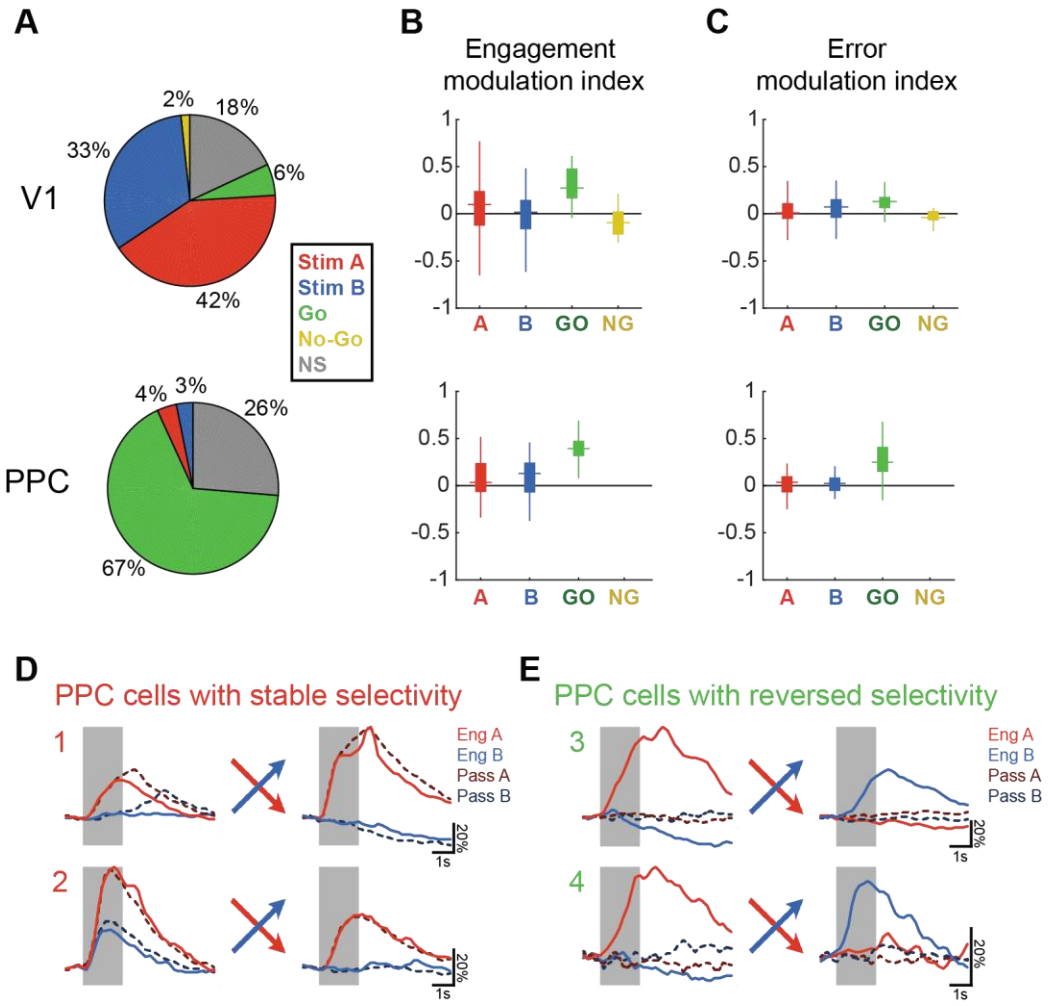

Figure 4.8. PPC cells with reversed selectivity exhibit stronger modulation by engagement and by error trials.

(**A**) Proportion of significantly responsive V1 (top) and PPC (bottom) neurons with stable stimulus selectivity (Stim A, 0°; or Stim B, 90°) or reversed selectivity (Go or No-Go). Cells with non-significant selectivity either before or after reversal are categorized as non-selective (NS). (**B**) Boxplots (showing median, interquartile range, and range) of engagement modulation index for each response group, computed as the mean before and after reversal. “Go”-selective neurons (green) show stronger engagement modulation, i.e. weaker passive responses, than stimulus-selective neurons (red or blue). (**C**) Boxplots of error modulation index for each response group, computed by comparing False Alarm trials with Correct Reject trials both before and after reversal. “Go”-selective PPC neurons show stronger error modulation (green) than stimulus-selective neurons (red or blue). (**D**) Responses of two PPC neurons with stable selectivity to Stimulus A both before and after reversal of reward contingency. Colors indicate stimulus (red, A; blue, B) and dark, dashed lines indicate passive trials. These neurons show a robust passive response in both conditions, and have a low engagement modulation index. (**E**) Same as (D) but for two PPC neurons with reversed selectivity after reversal of reward contingency. These neurons show virtually no passive response in both conditions, and have a high engagement modulation index.

Chapter 5: Dynamics of population encoding and decoding in parietal cortex during a sensorimotor decision.

5.1 Summary

Single neurons in higher cognitive regions like the posterior parietal cortex (PPC) can be widely variable in both the selectivity and temporal dynamics of their responses. Such heterogeneity of task-related signals is challenging to interpret at the level of individual responses, and may be better understood in the context of a distributed population-level code. Using a combination of statistical approaches and decoding analyses, we examined how populations of neurons in mouse PPC and primary visual cortex (V1) encoded information during performance of a visual discrimination task, and during passive viewing of the same stimuli. Although linear models using task-related signals could predict single-neuron calcium responses to some degree, model performance was greatly improved by the inclusion of inter-neuronal correlations, particularly for PPC. We thus moved to a population-level analysis of information coding in V1 and PPC. Using a time-dependent decoding approach, we find that task-specific information is encoded dynamically in PPC but not in V1, with different patterns of activity coding for the stimulus at different times within a trial. Mouse PPC therefore differs from V1 not only in its greater heterogeneity of individual responses, but also in the dynamic nature of its population code.

5.2 Introduction

Great progress has been made over the past decades of systems neuroscience by relating the responses of single neurons to externally measurable or controlled variables. This approach has been especially fruitful for the early sensory system, where individual neurons can express strong selectivity to sensory parameters such as orientation (Hubel and Wiesel, 1962), direction of motion (Britten et al., 1996), or disparity (Nienborg and Cumming, 2006). However, in higher

cognitive areas, such as the prefrontal cortex or the posterior parietal cortex (PPC), relating single neuron activity to individual task components is more challenging. Although some neurons in these areas may be tuned to a single task parameter, many others (and perhaps the majority) express mixed selectivity to multiple task components, and often with heterogeneous time-courses (Jun et al., 2010; Mante et al., 2013; Meister et al., 2013; Raposo et al., 2014; Rigotti et al., 2013). For example, during perceptual decision-making tasks, the activity of some PPC neurons appears to directly encode the accumulated evidence for a particular choice (Hanks et al., 2015; Roitman and Shadlen, 2002), leading some to suggest a primary role of PPC in the process of mapping sensory information onto appropriate motor actions (Andersen and Cui, 2009; Gold and Shadlen, 2007). However, closer examination of responsive PPC neurons during these decision tasks reveals a great deal of heterogeneity in both selectivity (Raposo et al., 2014; Rishel et al., 2013) and dynamics (Goard et al., 2016; Harvey et al., 2012; Morcos and Harvey, 2016). Individual neurons in PPC can simultaneously but separately encode both sensory accumulation and action selection signals (Bennur and Gold, 2011), and even decision-irrelevant visual signals can be mixed in with decision-relevant ones (Meister et al., 2013). Reading out task-relevant signals from a distributed and multiplexed population code would therefore require de-mixing them from other signals (Murakami and Mainen, 2015; Park et al., 2014).

Although there are many approaches to de-mixing various task signals from neuronal responses (Kobak et al., 2016; Machens et al., 2010; Mante et al., 2013), one of the simplest is to directly relate them to neural activity via an encoding model. Generalized linear models have been used to describe information encoding in early sensory areas (Pillow et al., 2005; Pillow et al., 2008), motor areas (Paninski et al., 2004; Truccolo et al., 2005) as well as higher cognitive regions (Barbieri et al., 2005; Park et al., 2014; Rorie et al., 2010), whether using recorded spiking data or surrogates such as imaged calcium responses (Chen et al., 2016; Pinto and Dan, 2015). These models have shown some success in dissociating the contributions of different task-related signals to neuronal activity (Park et al., 2014), even for areas like PPC which may encode more

elusive cognitive signals like attention (Bisley and Goldberg, 2010) or choice (Gold and Shadlen, 2007).

Understanding how neurons encode diverse sets of information is further complicated by the knowledge that the activity of cells in a network is often strongly coupled (Averbeck and Lee, 2006; Zohary et al., 1994). Trial-to-trial fluctuations in activity are often highly correlated between neurons, and these dependencies, often termed noise correlations, are usually seen as detrimental to population coding (Averbeck et al., 2006; Cohen and Kohn, 2011; Ecker et al., 2011; Zohary et al., 1994), but see (Kohn et al., 2016; Moreno-Bote et al., 2014). Although encoding models of early sensory (Pillow et al., 2008) and motor areas (Truccolo et al., 2005) have been built to incorporate inter-neuronal correlations, such models are rare for higher cognitive areas in the context of behavioral tasks.

In this work, we developed encoding models using population imaging data from mice performing a visual decision task. Heterogeneous neuronal responses were recorded from V1 and PPC during engaged task performance as well as during passive viewing of the same stimuli. Although linear models using task-related variables were successful in describing V1 activity to some degree, they performed more poorly for PPC. Measurements of noise correlations suggested that stronger correlations in PPC may have led to the poor prediction performance. By then extending the models to incorporate network fluctuations, we greatly improved our ability to explain single neuron activity, particularly for PPC.

In addition to an encoding perspective, which seeks to explain the responses of neurons given measured variables, one must also ask the complementary question of decoding: how might downstream neurons read-out task-relevant information from a population using biologically plausible mechanisms? This question is not trivial, given the heterogeneity and time-varying dynamics of responses in areas like PPC. Decoding of information may be relatively straightforward if the seemingly complex population responses are actually low-dimensional in nature (Ganguli et al., 2008; Murray et al., 2017). In this scenario, stable coding of information

allows for relatively straightforward readout using a set of fixed decoding weights. However, in some situations, coding in higher cognitive areas has been shown to be dynamic (Crowe et al., 2010; Harvey et al., 2012; Meyers et al., 2008; Stokes et al., 2013), with different patterns of neurons encoding the same information at different times. Readout in this scenario may be more challenging, as it may require time-varying weights which would be difficult to implement biologically. Whether such dynamic, high-dimensional codes are specific to certain brain areas or to particular types of tasks remains unclear.

To assess how information is encoded in V1 and PPC during a visual discrimination task, we employed a population decoding approach. We assessed whether V1 and PPC encoded information about the stimulus in a static or dynamic manner using a cross-temporal decoding approach. We find that task-specific information was encoded dynamically in PPC but not in V1, and that the dynamic code was present only in the context of task performance.

5.3 Experimental Procedures

5.3.1 Behavioral tasks, imaging data, and basic analyses

The behavioral tasks, imaging experiments, cell selection and preprocessing steps used to generate the experimental data used in this chapter are described in detail in Section 4.3. Briefly, mice were trained to perform a head-fixed, go/no-go visual discrimination task. Mice were trained to lick after presentation (2 s) of a target stimulus (horizontal grating drifting 0º from vertical) for reward (Hit), and to suppress licking (Correct Reject, CR) after a non-target (grating drifting 90º from vertical) stimulus. Incorrect licks to the non-target stimulus were punished (False Alarm, FA), whereas failure to lick to the target stimulus was not punished (Miss). A retractable lick spout was used to restrict licking responses to the period after stimulus offset (1.5 s) during Engaged trials, and was not presented at all to the animal during blocks of Passive trials. In a subset of imaging

sessions, the contrast of the stimuli was varied randomly from trial to trial between one of six values (2, 4, 8, 16, 32, or 64%). Those sessions are included only in some of the decoding analyses (**Figure 5.4E,F**; **Figure 5.6B,C**; **Figure 5.7C,F**). Finally, in a separate set of imaging sessions, the delay between stimulus offset and spout presentation was varied randomly from trial to trial between 0, 3, and 6 seconds (see Section 2.3). Stimulus contrast was held constant in these sessions. This data was analyzed only in one of the decoding analyses (**Figure 5.6D,E**).

Calcium imaging was performed on large populations of neurons from either V1 or PPC using volumetric imaging (5 Hz). Blocks of engaged task performance were interleaved with blocks of passive viewing of the same sequence of stimuli. After image preprocessing and $\Delta F/F$ extraction, traces were sorted by trial type and condition. A neuron was considered task responsive and included for further analysis if its mean $\Delta F/F$ during the last 1.6 s (8 frames) of the stimulus period was significantly ($p < 0.01$, t-test) greater than the pre-stimulus baseline (1 s), for either hit or correct reject trials. Neurons also had to meet a signal-to-noise criterion, needing a trial-averaged response that exceeded a threshold of at least 2 standard deviations above baseline during either the stimulus or choice period. For data acquired during the variable contrast task, neurons were considered significantly responsive if the mean $\Delta F/F$ during the stimulus period was significantly above threshold for at least two of the six contrasts of the same stimulus (target or non-target). Neurons were marked as target- or non-target-preferring based on their mean response during Engaged trials.

Response latencies for each neuron (**Figure 5.5A**) were estimated by first up-sampling trial-averaged responses to 500 Hz. The latency was computed as the time at which a neuron's response exceeded 1 standard deviation above the baseline distribution. Responses that did not reach the threshold of 2 standard deviations were not included in latency distributions. Histograms were plotted using the latency of each neuron in response to its preferred stimulus.

5.3.2 Signal and noise correlations

Signal correlations between pairs of simultaneously recorded neurons were calculated as the Pearson CC between trial-averaged responses. Noise correlations were computed by first subtracting the trial average from single trial responses and then computing the Pearson CC between the mean-subtracted responses (Hofer et al., 2011). Correlations were computed separately for Engaged and Passive conditions, and using only correct trials for the Engaged condition. All cumulative histograms (**Figure 5.2C, F**) were computed by finding the median correlation for each cell with all other simultaneously recorded cells with the same stimulus preference (measured during Engaged trials), and then pooling across all imaging sessions. The distance dependence of noise correlations was computed by binning all pairs of simultaneously recorded neurons (with the same stimulus preference) by their Euclidean distance in 20 μm bins up to 300 μm, and then taking the average within each bin.

5.3.3 Generalized linear model: using task predictors only

We used a generalized linear model (GLM) to regress recorded calcium signals against a time series of task events (Chen et al., 2016; Miri et al., 2011; Pinto and Dan, 2015). Because spike inference was not performed and a linear relationship between calcium activity and the predictors was assumed, this modeling approach is equivalent to multiple linear regression with lagged predictors.

Calcium responses for each cell were Z-scored and modeled as the linear combination of various task events, each convolved with a filter:

$$\hat{y}_t = \beta_0 + \sum_{i=0}^{15} \beta_i^{S\_tar} x_{t-i}^{S\_tar} + \sum_{i=0}^{15} \beta_i^{S\_nt} x_{t-i}^{S\_nt} + \beta_0^{E} x^{E} + \sum_{i=0}^{15} \beta_i^{E\_tar} x_{t-i}^{E\_tar} + \sum_{i=0}^{15} \beta_i^{E\_nt} x_{t-i}^{E\_nt} + \sum_{i=-7}^{15} \beta_i^{L} x_{t-i}^{L}$$

The response of a neuron at frame t is modelled ($\hat{y}_t$) by the sum of a bias term (β_0) and the weighted (β_i) sum of various additional binary predictors at different lags (i). Binary predictors for the target stimulus ($x_t^{S\_tar}$) and non-target stimulus ($x_t^{S\_nt}$) indicated the duration of stimulus presentation in either engaged or passive trials. Binary predictors for engagement included a constant offset (x^E) that was 1 during engaged trials and 0 otherwise, as well as stimulus predictors ($x_t^{E\_tar}$, $x_t^{E\_nt}$) that indicated the duration of stimulus presentation during engaged trials only. Binary predictors (x_t^L) for licking indicated the duration of lick bouts, which were defined as groups of licks with an inter-lick interval less than 1 second. The number of lags (15 samples = 3 seconds) were chosen to capture the full dynamics of the calcium response (GCaMP6s has slow offset kinetics (Chen et al., 2013)). Lags were chosen to be strictly positive (causal) for stimulus and engagement predictors, but both positive and negative (anti-causal) for licking predictors. The final model had 88 coefficients including a constant bias term.

Models were fit using ridge regression using procedures similar to that of previous studies (Chen et al., 2016; Huth et al., 2012; Pinto and Dan, 2015). We first set aside 20% of trials from each condition (Hit, Correct Reject, Miss, False Alarm, Passive Target, Passive Nontarget) for testing. A regularization parameter λ was estimated for each cell from among a range of λ values (10^{-2} to 10^4) using cross-validation. Model performance was measured by computing the correlation coefficient (CC) between predicted and measured activity. Five-fold cross-validation was used to choose the λ that maximized the performance on the remaining 80% of training trials. The final model for each cell was fit using the best λ, and then performance was evaluated by measuring CC of predictions for the holdout test set. Statistical significance of the predictions was evaluated by bootstrapping the test data set 1000 times and obtaining a distribution of CC values. A p-value was calculated as the proportion of iterations with CC ≤ 0, and neurons with $p < 0.05$ were considered to have significant model fits.

5.3.4 Generalized linear model: using task and network predictors

Before explaining the combined task and network model, we will first switch to a simplified matrix description of the task-only model laid out in Section 5.3.3 for clarity. The task-only GLM models the responses of a population of n neurons across t time bins with p predictors as follows:

$$Y = XB + R$$

Where Y is a $t \times n$ matrix of z-scored calcium responses, X is a $t \times p$ matrix of predictors, B is a $p \times n$ matrix of task-related coefficients and R is a $t \times n$ matrix of residuals. The model prediction for the task-only GLM is simply $\hat{Y}_1 = XB$.

In order to incorporate network predictors in the GLM, we considered two possible approaches. A simultaneous approach (Pillow et al., 2008; Truccolo et al., 2005) would incorporate the network terms as an additional term in the above equation:

$$Y = XB + YW + R$$

Where W is a $n \times n$ matrix of inter-neuronal network coupling weights, with the diagonal of W constrained to be equal to 0. In such a model, weights for task and network predictors would directly compete against each other. Given the presence of high signal correlations between neurons (**Figure 5.2C**), this model would dramatically reduce the task weights, and also make interpretation of network weights difficult.

An alternative approach is a sequential model (Malik et al., 2011), in which network weights were estimated using the residuals of the task model:

$$Y = XB + R$$

$$R = RW + E$$

Where W is again a $n \times n$ matrix of inter-neuronal network coupling weights with a zero diagonal, and E is the remaining error. The final prediction can thus be computed by summing the outputs of the two models:

$$\hat{Y}_1 = XB$$

$$\hat{R} = RW = \left(Y - \hat{Y}_1\right)W$$

$$\hat{Y}_2 = \hat{Y}_1 + \hat{R} = \hat{Y}_1 + \left(Y - \hat{Y}_1\right)W$$

In this sequential model, which is depicted and evaluated in **Figure 5.3**, network weights can be naturally related to noise correlations, as both represent residual trial-to-trial fluctuations after "signal" components have been removed.

To fit the "task + network" model, the same training trials (consisting 80% of the data) were used as for the "task-only" model (Section 5.3.3). Residuals for each neuron were computed by subtracting model predictions from the training data. We z-scored the residuals by dividing by the standard deviation to normalize all network weights, but kept the normalization factor for each cell in memory. For each cell, the simultaneous (lag = 0) residual activity of all other cells in the network (R) were used as predictors. We avoided using additional lags to keep model complexity low, which we considered reasonable given the low time-resolution of calcium responses (200 ms). Models were fit using ridge regression, as with the task-only model (Section 5.3.3). We computed a single regularization coefficient λ for each population of cells, estimated using a random 20% subset of the cells to save computation time. Prediction performance was not significantly improved using individual values of λ for each cell for a handful of test populations (data not shown). As before, λ was estimated using five-fold cross-validation and prediction performance was evaluated using the CC between predicted and measured responses in the training dataset. The λ that maximized the average CC across the random subset of cells used for cross-validation was then used to fit the network model for all cells.

Final prediction performance was evaluated on the test dataset. Residuals were generated by subtracting the predictions of the task-only model from the test data ($R = Y - \hat{Y}_1$). The residuals of a given cell were then predicted using the network model and the residuals of other cells in the network ($\hat{R} = RW$). The predicted residuals for each cell were multiplied by the normalization constant and the summed with the output of the task-only model ($\hat{Y}_1$) to generate final model prediction ($\hat{Y}_2 = \hat{Y}_1 + \hat{R}$). As described in Section 5.3.3, statistical significance of predictions was evaluated using a bootstrap test (1000 iterations).

To assess the relationship of model performance to noise correlations (**Figure 5.3E**), we binned neurons based on their median absolute noise correlation with other cells (bin size 0.05) and averaged the model performance within each bin. We also trained models using subsets of the network inputs (**Figure 5.3F**). For each cell, we sorted all other cells (which we refer to as "network inputs" for clarity) based on their absolute noise correlation with the modelled cell, in ascending order. The 20, 40, 60, and 80% of inputs with the weakest noise correlation strength were each used to train a separate model of the cell's residual response (R). Model performance for each of these sub-network models was evaluated by summing with the task-only model predictions, as with the full-network model.

5.3.5 General decoding analyses

All decoders were linear support vector machines (SVM) (Cortes and Vapnik, 1995) implemented using the libsvm library in MATLAB. We chose SVM because it is a state-of-the-art machine learning classifier that yields high performance while minimizing overfitting to training data, and has been shown to perform well on neural population data (Graf et al., 2011). A linear kernel was chosen because its architecture could be easily implemented using biologically plausible mechanisms as a thresholded sum of weighted synaptic weights.

Data was preprocessed by first z-scoring the data for each cell, sorting into trials, and then smoothing single-trial responses across 3 consecutive time bins (600 ms window) to reduce noise. For display purposes, all performance curves are plotted with respect to the last time bin in the 600ms window. Performance was evaluated using at least 10 iterations of five-fold cross-validation. On each iteration, the number of trials in each condition (e.g.: Target versus Nontarget trials) were equalized by pseudo-random subsampling from the condition with the larger number of trials. The subsampling procedure was pseudo-random in that all individual trials were selected evenly, such that the number of iterations per trial were roughly equal. This was particularly important in the case of error trial decoding (**Figure 5.7**) because some sessions had as few as five error trials. We increased the number of iterations beyond 10 (up to 25) to ensure that every trial of the subsampled condition was used in at least one iteration.

We used a C-support vector classifier, and initially tested the performance of various values of the regularization variable C on a few test datasets using cross-validation. We found that a value of C = 10 performed optimally across these datasets, and used this optimized value for all other datasets.

In each iteration of five-fold cross-validation, decoders were trained using the population activity at a single time bin using 80% of the trials and then tested on the remaining 20% of trials at the same time bin. Performance was averaged across the five folds, and then across the 10+ iterations. Significance of performance was assessed by shuffling trial labels 1000 times and generating a distribution of accuracies given a random 50/50 guess. Significant ($p<0.05$) decoding was achieved if the true performance exceeded the range spanned by the center 95% of values from this shuffled distribution. For decoding of Target versus Non-target (**Figure 5.4**, **Figure 5.5**), in which there were many trials available for decoding, this significance threshold was around 58%. For decoding with error trials (**Figure 5.7**), there were less trials available, resulting in a significance threshold of about 66% for False Alarm trials and 74% for Miss trials.

To assess decoding performance with subsets of neurons during Engaged or Passive conditions (**Figure 5.4C, D**), we first sorted neurons in ascending order (Morcos and Harvey, 2016) based on their absolute selectivity index, which was directly related to the area under the ROC curve (see Figure 4.2G, Section 4.3.6). The selectivity index was computed based on responses during the latter part of the stimulus period (0.8 to 2 s). Neurons with the lowest magnitude selectivity were successively added in 20% percentiles (20%, 40%, 60%, 80%, 100%), and five separate decoders were trained for each time bin, and each condition. Sorting was done separately for the two conditions (Engaged, Passive) based on individual selectivity during each condition. Average performance for each subset and condition (**Figure 5.4D**) was computed by averaging across the latter part of the stimulus period (0.8 to 2 s) and across populations.

To assess decoding performance across various contrasts (**Figure 5.4E; Figure 5.7C, F**), decoders were trained using all cells in a population and were not given information about contrast. Performance was assessed separately for trials of each contrast level, and average performance for each contrast and condition (**Figure 5.4F**) was computed by averaging across the latter part of the stimulus period (0.8 to 2 s) and across populations.

5.3.6 Time-dependent decoding analyses

To assess time-dependent coding (**Figure 5.5**, **Figure 5.6**), we used a cross-temporal decoding approach (Crowe et al., 2010; Meyers et al., 2008; Stokes et al., 2013). In addition to training and testing decoders at the same time bin ("Within" time bin decoding), we also trained decoders at various single time bins and then tested their performance across all time bins ("Across" time bin decoding). Specifically, we trained 12 different "Across" time bin decoders per dataset and condition, each trained at a different, non-overlapping 600 ms time window (centered at -0.6, 0, 0.6, 1.2, etc., up to 6.0 s relative to stimulus onset), and then tested the performance of each on

all time bins. Performance was evaluated using 10+ iterations of five-fold cross-validation as described above. Therefore if a given trial was used for training a model, it was never re-used for testing that model, even if the data for training and testing came from different time bins.

To test whether a stable readout could be used even in the context of a dynamic code (**Figure 5.6A,B**), we trained a decoder to discriminate Target from Non-target trials using data from multiple time-bins ("Multiple" time bin decoding). Specifically, population activity from six different non-overlapping 600ms time windows (centered at 0.6, 1.2, 1.8, 2.4, 3.0, and 3.6 s relative to stimulus onset) was taken from a subset of trials to train the decoder, with data from each time bin used as separate training trials. A single Target trial, for example, would contribute six different population patterns for the decoder to contrast with patterns from Non-target trials. Performance was evaluated using ten iterations of five-fold cross-validation as described above. This approach differs from "Within" time bin decoding because it represents a single weight vector that is tested across different time bins. It also differs from the "Across" time bin decoding approach in that the decoder is not optimized for one particular time window.

Time-dependent decoding was also performed on data acquired during the variable contrast task (**Figure 5.6B,C**), or during the variable delay task (**Figure 5.6D,E**). Choice decoders were trained on the variable contrast data to distinguish Hit versus Correct Reject trials irrespective of contrast. Contrast decoders were trained to distinguish High contrast Hit trials (32% and 64%) from Low contrast Hit trials (2% and 4%). Decoding in the variable delay task was performed separately for each delay (0, 3, and 6 s) by comparing Hit versus Correct Reject trials. A total of 22 separate "Across" time bin decoders were trained for the 6 s delay trials, each using a different non-overlapping 600 ms time window.

5.4 Results

5.4.1 Relating V1 and PPC responses to task events using a generalized linear model

As previously described in Chapter 4, calcium responses were obtained from populations of neurons in mouse V1 and PPC during a head-fixed, go/no-go visual discrimination task using two-photon calcium imaging. In this task, mice were trained to lick after presentation (2 s) of a target stimulus (horizontal grating drifting 0º from vertical) for reward (Hit), and to suppress licking (Correct Reject, CR) after a non-target (grating drifting 90º from vertical) stimulus. Incorrect licks to the non-target stimulus were punished (False Alarm, FA), whereas failure to lick to the target stimulus was not punished (Miss). Blocks of engaged task performance (Engaged) were interleaved with blocks of passive viewing of the same sequence of stimuli (Passive). A retractable lick spout was used to restrict licking responses to the period after stimulus offset during Engaged trials, and was not presented at all to the animal to signal blocks of Passive trials.

We first sought to determine the degree to which single neuron responses could be explained by a linear combination of task parameters. To disambiguate the contributions of various task events to V1 and PPC responses, we used a well-established paradigm, the generalized linear model (GLM), which has been used to model both spiking data (Park et al., 2014; Pillow et al., 2008) and calcium responses (Chen et al., 2016; Miri et al., 2011; Pinto and Dan, 2015). Z-scored calcium responses of each cell were fit to a linear combination of ongoing task events, including stimulus presentation, task engagement, and licking (**Figure 5.1A**). We simultaneously fit the responses of cells during Engaged and Passive trials, and then measured the performance of the model on a separate test dataset using the correlation coefficient (CC) between model predictions and the actual data (**Figure 5.1B**). Using neurons previously identified as significantly responsive (see Section 5.3.1), the average CC was higher for V1 (0.54 ± 0.01, mean ± SEM, n = 1915 cells) than for PPC (0.42 ± 0.01, n = 3524 cells), corresponding to an average explained variance of 33% for V1 versus 20% for PPC. V1 neurons are therefore better described by a linear model based on task-related events.

We illustrate the performance of the "task-only" model, named thus because it uses only task events as predictors, on two example neurons from each area (**Figure 5.1C-F**). Trial-averaged responses for each trial type (Hit, CR, Miss, FA, Passive Target, and Passive Nontarget) were decomposed into three components: a Stimulus component, an Engagement component, and a Licking component. Both example V1 cells had a large Stimulus component, as they responded similarly to stimuli in both Engaged and Passive conditions (**Figure 5.1C**). In contrast, both example PPC cells had very weak Passive responses, and therefore weak Stimulus components (**Figure 5.1E**).

Responses on Engaged trials incorporated the other two model components. The second component, task engagement, was modeled using the sum of a constant offset and a stimulus-locked signal that exclusively occurred on Engaged trials. The third component, licking, also included lags that were anti-causal, such that calcium responses that preceded or followed licking bouts were incorporated. One V1 neuron (see Cell 1 in **Figure 5.1C**) exhibited weak modulation of their responses with engagement, but both V1 neurons had weak licking components, as neither neuron exhibited much activity on False Alarm trials. In contrast, both PPC neurons had a strong licking component (**Figure 5.1E**), which predicted the increase of activity on FA trials. In the case of one PPC neuron (see Cell 4 in **Figure 5.1E**), licking alone was not sufficient to describe the enhancement of response on Hit trials, and therefore a target-selective engagement component is included. The models for all four example cells have relatively high CC on holdout test data trials (**Figure 5.1D, F**). These examples demonstrate face validity for the model's ability to separate different task-related components to responses in V1 and PPC neurons.

We quantified the strength of each of these three components by measuring the median absolute weight value, which is normalized and thus can be directly compared across cells (**Figure 5.1G**). On average, V1 neurons have much stronger stimulus weights than PPC neurons (V1: 0.067 ± 0.001; PPC: 0.027 ± 0.001; mean ± SEM), whereas PPC neurons have slightly stronger engagement (V1: 0.058 ± 0.001; PPC: 0.068 ± 0.001) and much stronger licking weights

(V1: 0.057 ± 0.001; PPC: 0.101 ± 0.001). The qualitative differences in model weights between areas match the conclusions from previous single neuron analyses of responses in passive conditions (see Figure 4.2) and on error trials (see Figure 4.3), lending further validity to the model's ability to dissociate the contributions of stimulus, engagement, and licking.

5.4.2 Noise correlations are stronger in PPC than in V1

The task-only model was sufficient to predict response properties of neurons from both areas, but a large proportion of variance was left unexplained, particularly for PPC neurons (20% variance explained). One possible source of unexplained variance is the fluctuation of network activity from trial-to-trial. These inter-neuronal dependencies, often termed noise correlations, can be relatively strong in sensory cortex, especially during inattentive states (Cohen and Kohn, 2011; Cohen and Maunsell, 2009; Zohary et al., 1994). Before incorporating them in our model, we first characterized the strength of noise correlations in V1 and PPC during task performance and passive viewing.

We first separated each population based on their stimulus preference during Engaged trials. Neurons with the same stimulus preference naturally exhibited large signal correlations (computed as the correlation between trial-averaged responses during correct Engaged trials) in both V1 (**Figure 5.2A**) and PPC (**Figure 5.2B**), though there was a much larger proportion of Target-preferring neurons in PPC (see Figure 4.2E). In both areas, signal correlations between neurons with the same stimulus preference increased with engagement (V1, passive: 0.65 ± 0.01; engaged: 0.78 ± 0.01; PPC, passive: 0.45 ± 0.01; engaged: 0.81 ± 0.01; **Figure 5.2C**), especially in PPC which exhibits very weak passive responses (see Figure 4.2C,F).

We then computed noise correlations in both Engaged and Passive conditions by correlating trial-to-trial fluctuations across simultaneously recorded neurons. Noise correlations between neurons with the same stimulus preference tended to be stronger than noise correlations

between pairs with opposite stimulus preferences (data not shown). Focusing on pairs with the same stimulus preference, V1 neurons exhibited similar strength of noise correlations in Engaged versus Passive conditions (passive: 0.11 ± 0.01; engaged: 0.11 ± 0.01; **Figure 5.2D**). In contrast, PPC cells tended to have stronger noise correlations in Passive versus Engaged conditions (passive: 0.25 ± 0.01; engaged: 0.19 ± 0.01; **Figure 5.2E**). Comparing across all cells and all populations, noise correlations were stronger in PPC compared to V1, especially during Passive trials (**Figure 5.2F**). These correlations depended on the anatomical distance between neurons (**Figure 5.2G**) (Rikhye and Sur, 2015). These results indicate that network correlations strongly affect neural responses, particularly in PPC.

5.4.3 Addition of network coupling terms in GLM improves model prediction

We then sought to incorporate correlations between neurons in our GLM as network coupling terms to see if they would improve prediction performance (Pillow et al., 2008; Truccolo et al., 2005). We reasoned that because of large signal correlations between neurons (Figure 5.2A-C), directly incorporating the network terms into the original model would mask the contribution of the various task events. Therefore, we first fit the data to a task-only model which captured the "signal" components (Malik et al., 2011), and then fit the residual errors for each neuron using a network model (**Figure 5.3A**). An added advantage of this framework is that it allows direct comparison of network weights with noise correlations, as they are similarly computed by first subtracting off the trial-averaged "signal". Prediction performance using this sequential fitting procedure was similar to results using simultaneous fitting of task and network terms (data not shown). To keep the number of regressors in the model low, we used only a single (0-frame) lag for each neuron in the network, which we considered reasonable given the low time resolution of the calcium responses (200 ms).

We measured the performance of the “task + network” GLM and compared it to the “task-only” model, again using the correlation coefficient of the model predictions with test data not used to train either model (**Figure 5.3B**). Prediction performance improved for cells in V1 (CC, task-only: 0.54 ± 0.01; task+net: 0.75 ± 0.01), but even more so for neurons in PPC (CC, task-only: 0.42 ± 0.01; task+net: 0.77 ± 0.01), leading to an average explained variance of 57% for V1 and 60% for PPC. The large improvement in PPC was expected given the high level of noise correlations measured in PPC responses. Model performance was higher for Passive trials compared to Engaged trials (data not shown), also predicted by the higher levels of noise correlations on Passive trials.

We illustrate the improved performance of the “task + network” model using a few single-trial responses from one V1 neuron (**Figure 5.3C**) and one PPC neuron (**Figure 5.3D**). Although the “task-only” GLM was proficient at predicting trial-averaged responses, it failed to capture variations in responses between trials with the same behavioral outcome, e.g. two different Hit trials. However, as seen in both example cells, incorporation of network coupling terms improved the model’s ability to predict responses on single trials.

We examined the relationship between model performance and the strength of noise correlations for a given cell (**Figure 5.3E**). There was a strong positive correlation for both V1 ($r = 0.21$, $p < 10^{-9}$) and PPC ($r = 0.42$, $p < 10^{-9}$), indicating that cells that were more strongly influenced by network inputs could be better explained using the “task + network” model. Interestingly, noise correlation strength was inversely correlated with task-only model performance for V1 ($r = -0.22$, $p < 10^{-9}$) but not for PPC ($r = 0.04$, $p > 0.01$). This suggests that for V1 neurons, the ability of a linear model to predict responses based on task parameters (i.e. stimulus) is corrupted by the presence of noise correlations (Zohary et al., 1994).

How much of the network is needed to improve model performance? It’s possible, for example, that the improved performance using network terms is due to a small fraction of highly correlated inputs. Indeed, for many PPC cells, a few inputs with the strongest noise correlation

were often given disproportionately high model weights (data not shown). To examine the relationship between model performance and network size, we first sorted the inputs for each modeled cell in ascending order based on the strength of its noise correlation with the cell (**Figure 5.3F**, left). We then trained new models for each cell, using only a subset of the network which contained the inputs with the weakest noise correlations (**Figure 5.3F**, right). Models using 80% of the cells (thus eliminating the 20% of inputs with the strongest noise correlations) performed nearly as well as the full network model for both V1 (CC: 0.69 versus 0.75 for full model) and PPC (CC: 0.71 versus 0.77 for full model). This suggests that the network inputs onto a cell are largely redundant, and specific inputs are not required in order to improve prediction performance. Indeed, model performance was dramatically improved for PPC even using just the 20% of network inputs with the weakest noise correlation (CC: 0.57 versus 0.42 for task-only model).

Overall these results demonstrate that both V1 and PPC responses can be better described when incorporating network correlations, indicating the importance of population-level descriptions of neural data.

5.4.4 Decoding stimulus from V1 and PPC during engaged and passive conditions

The large influence of network correlations on single neural responses led us to reevaluate the data from a population-level perspective. In previous analyses (see Chapter 4), we evaluated the information coding properties of V1 and PPC neurons, which could vary widely in their individual sensitivity to engagement, contrast, or contingency reversal. But can an observer read out the appropriate information given a whole population of heterogeneous neurons? To answer this question, we used a decoding approach to assess how information was coded in V1 and PPC population activity during engaged and passive conditions. We chose a machine learning classifier, the linear support vector machine (SVM) (Cortes and Vapnik, 1995), both because of

its robustness to overfitting and because its architecture could be realistically implemented by a downstream readout neuron using a weighted sum of synaptic weights.

We first built decoders that could classify Target versus Nontarget trials during either Engaged or Passive conditions, to test how stimulus information in V1 and PPC were affected by engagement (**Figure 5.4**). Separate decoders were trained for each time bin and each condition. If we used all cells to train the decoder, performance rapidly achieved ceiling levels for V1 in both Passive and Engaged conditions (**Figure 5.4A**). This is expected given high numbers of selective neurons in V1 during both conditions (see Figure 4.2G). PPC neurons typically had weak responses on Passive trials, and so we expected decoding performance to be correspondingly weak. However, PPC decoder performance was significant in both Engaged and Passive conditions (**Figure 5.4A**). Variability across individual populations was much greater in PPC than in V1, but across conditions there was significant performance even during Passive trials (**Figure 5.4B**).

Is information about the stimulus distributed across the population or can it be recovered using just a subset of neurons? To address this question, we first sorted neurons by their absolute selectivity index (see Figure 4.2G), and trained decoders with successively more and more neurons, starting with the 20% neurons with the weakest selectivity (**Figure 5.4C, D**). V1 neurons generally had strong selectivity and so even the 20-40% least selective V1 neurons could perform well above chance at predicting stimulus identity in both Engaged and Passive conditions. By contrast the 60% least selective PPC cells could not be used to predict the stimulus above chance levels in Passive conditions. This implies that passive coding of stimuli is limited to a small subset of neurons in PPC (see also Figure 4.2F). Decoding performance with a given subset was not significantly higher than performance using the best neuron in that subset (data not shown), in contrast to recent findings in mouse PPC during a navigational accumulation of evidence task (Morcos and Harvey, 2016). This suggests that in our task, information about stimulus within single neurons is strong enough that a distributed population code may not be necessary.

Could PPC populations encode stimulus differently under conditions of lower input strength? In some sessions we varied the stimulus contrasts from trial to trial, and so we trained decoders on these datasets to discriminate stimulus identity, independent of contrast (**Figure 5.4E, F**). We found that performance improved with contrast across areas and conditions, although V1 performed at near ceiling levels even at the very lowest contrasts (4-8%). Overall, these results indicate that information about stimulus is therefore encoded very reliably in V1 populations regardless of behavioral state. However, stimulus information in PPC populations is less reliable, particularly during Passive conditions.

5.4.5 Time-dependent decoding in PPC but not V1

The decoding approach described above revealed that stimulus information can be easily read-out from both V1 and PPC, especially during Engaged conditions. But is there any difference in the way these areas encode this information? Prior work examining PPC responses in the context of decision tasks have revealed that activity patterns can vary dramatically in time across neurons (Meister et al., 2013; Raposo et al., 2014), and in some cases even forming choice-specific sequences (Harvey et al., 2012; Morcos and Harvey, 2016). Indeed, our own prior work using a variable delay task (Goard et al., 2016) has shown that PPC neurons exhibit a wide range of heterogeneity in the temporal profile of their responses (see Figure 2.3).

One possibility is that information in PPC may be encoded dynamically, with different neural activity patterns coding for the same information at different times during the task. Such time-dependent coding could be revealed using a cross-temporal analysis (Crowe et al., 2010; Meyers et al., 2008; Stokes et al., 2013), in which a decoder is trained at one time bin but tested at different time bins. If information is encoded using a single, stationary pattern of activity, high decoding performance at one time bin should generalize when tested across a wide range of time bins. By contrast, if information is encoded dynamically, then high decoding accuracy when

trained at one time bin should be limited to just that time bin or to a narrow range of neighboring time bins.

One preliminary piece of evidence that PPC may differ from V1 in temporal coding of information comes from examining the distribution of response onset latencies in both areas (**Figure 5.5A**). V1 neurons fired within a narrow range of latencies relative to stimulus onset in both Engaged and Passive conditions. By contrast, PPC neurons exhibited a broad range of latencies spanning the full 2-second stimulus period. This suggests the possibility that information may be carried by different PPC neurons at different points in time during the trial.

We then trained decoders to discriminate Target and Nontarget trials using V1 or PPC population data from a single time bin, and then tested performance across all other time bins (**Figure 5.5B-F**, "Across" time bin decoders). We trained separate decoders for Engaged and Passive conditions, and compared performance with decoders trained and tested within the same time bin (**Figure 5.5B**, "Within" time bin decoders, same as Figure 5.4A). As predicted, training and testing within the same time bin provided a ceiling level of performance for each condition. For V1, decoders trained at a single time bin and tested across different time bins performed very similarly, with high levels of accuracy across a wide range of time bins in both Engaged and Passive conditions. The average performance across populations is also plotted in a heatmap in which the different rows represent the training times of different decoders (**Figure 5.5C**). The region of high accuracy for V1 is large and does not shift much with training time.

By contrast, for PPC, "across" time bin decoders exhibited a narrower range of high decoding performance when trained on Engaged trials, with the peak in accuracy shifting forward in time as different time bins were used for training the decoder (**Figure 5.5B**). This can be seen in the heatmap (**Figure 5.5C**) as a relatively narrow band of high accuracy that follows the diagonal of the matrix. Dynamic coding was limited to the Engaged condition, as decoders trained on Passive trials at different time bins exhibited similar performance curves that did not shift with training time.

To quantify the dynamic or static nature of the decoders and its variability across individual populations, we measured different features of the decoding performance curves for each brain area and condition, as a function of training time (**Figure 5.5D-F**). As shown previously (see Figure 5.4A-B), the peak performance of the decoder was higher in V1 compared to PPC, and higher during Engaged trials compared to Passive (**Figure 5.5D**). Of note is that performance levels were relatively low when using the first few time bins within PPC. Low performance levels introduce noise in the following two measurements and thus these time bins should be interpreted with caution.

We then measured the width of the decoding peak for each curve as the time above a threshold performance of 75% (**Figure 5.5E**). Decoding width was consistently larger in V1 than in PPC, in both Engaged and Passive conditions. However, smaller decoding widths alone do not indicate dynamic coding, as they may simply reflect weaker overall performance, as in the case of PPC during Passive trials. We therefore also measured the time of the decoding peak as a function of training time, measured as the center-of-mass of the decoding curve (**Figure 5.5F**). The center-of-mass was largely stable for V1 across conditions, and for PPC during Passive conditions. By contrast, the center-of-mass for PPC during Engaged trials was highly dependent on training time. The narrower decoding width (**Figure 5.5E**) and shifting decoding peak (**Figure 5.5F**) suggests that dynamic coding is specific to PPC during Engaged conditions. PPC therefore exhibits a clear difference from V1 in that stimulus information is encoded in dynamic activity patterns, specifically during task performance.

5.4.6 Time-dependent coding during variable contrast and variable delay tasks

To further probe the nature of time-dependent coding in PPC, we performed three additional analyses. First we asked whether the dynamic activity patterns in PPC necessitated a dynamic readout mechanism. Others have demonstrated in prefrontal cortex that an apparently dynamic

code can coexist with stable readout (Murray et al., 2017). This could occur, for instance, if the population trajectories are distinct enough that a single linear hyperplane can separate them in neural state-space (Buonomano and Maass, 2009). While our cross-temporal approach trained decoders to be optimal at a single time-bin, it's possible that a decoder trained across multiple time bins would achieve high performance that is stable with time, even though it may be sub-optimal for any specific time point.

To test the possibility of stable readout, we trained decoders to discriminate Target from Non-target trials using population activity sampled at multiple distinct time-bins (**Figure 5.6A**). Importantly, individual time-bins were counted as separate instances for training, and performance was evaluated using cross-validation (see Section 5.3.6). This "Multiple" time-bin decoder performed nearly as well as "Within" time-bin decoding across time, with a very broad performance width, even though it utilized weights that did not vary with time. This was true for both areas and conditions, including in PPC during Engaged conditions, where "Across" time-bin decoders had relatively narrow widths. This argues that the presence of dynamic activity patterns can coexist with stable readout mechanisms.

In a second analysis, we took advantage of the variable contrast task to see whether signals reflecting stimulus contrast were also encoded in a dynamic manner (**Figure 5.6B,C**). We trained separate decoders to discriminate choice (Hit versus Correct Reject) and contrast (High contrast Hit versus Low contrast Hit) using the same V1 and PPC populations. Choice decoders were agnostic to stimulus contrast, and behaved similarly to the single contrast data. Specifically, coding of choice appeared to be dynamic in PPC but not in V1, nor in passive conditions (data not shown). By contrast, coding of contrast appeared to be stable in both PPC and in V1, although overall decoding performance was weaker in PPC. While more work needs to be done to determine the robustness of this phenomenon, these results suggest that the same PPC

population may be able to dynamically encode one variable, while stably encoding another (Crowe et al., 2010).

Lastly, the presence of dynamic coding in PPC in this task may appear to conflict with our previous results using a variable delay task, where PPC neurons exhibited stable activity during the delay period (see Figures 2.3, 2.5). We thus re-analyzed data from the variable delay task using the same cross-temporal analyses to decode Hit versus Correct Reject trials across the three delays (**Figure 5.6D,E**). Decoders trained on the 0 s delay trials behaved similarly to those trained on the single delay data (see Figure 5.5B), with more dynamism in PPC than in V1. However, increasing the delay between stimulus and response revealed a more complicated picture in both areas. Decoders trained in the stimulus epoch behaved very differently from those trained in the response epoch for both PPC and V1. This finding is consistent with the presence of response-period neurons in V1 (Figure 2.3) that are easier to identify in the longer delay task. In addition, decoders trained in the delay period exhibited large decoding widths for both areas, arguing against dynamic changes in activity during the delay. While more work needs to be done to characterize these findings, this preliminary analysis makes clear that dynamic coding is not present during the delay period of the task in either PPC or V1. Instead it is possible that the more rapid evolution of responses during the stimulus and response epochs accounts for the dynamic coding observed during 0 s delay trials.

5.4.7 Decoding stimulus and choice using error trials

What is the nature of the information carried by V1 and PPC populations during the task? The above decoding approach was limited to correctly performed trials, in which animals both observed a different stimulus (Target or Nontarget) and made a different choice (Lick or No-lick). Thus, the above decoding accuracy curves could represent either stimulus or choice information, or a combination of two separate signals, potentially with differing time-courses.

We had previously analyzed whether single neurons were sensitive to stimulus or choice using error trials to disambiguate stimulus and choice (see Figure 4.4, Figure 4.6). We now take a population decoding approach to see whether such information may be distributed across neurons. We again took V1 and PPC responses from behavioral sessions in which the mouse committed at least five misses or five false alarms (FA). We then trained population decoders to separately distinguish FA trials from Hit trials to assess stimulus coding, and from Correct Reject (CR) trials to assess choice coding (**Figure 5.7A**). All error trial decoders were trained and tested using the same time bins, and the performance across all behavioral sessions with sufficient numbers of error trials was averaged together. We performed this analysis using sessions with a single contrast (**Figure 5.7B**) as well as sessions with variable contrast (**Figure 5.7C**), keeping decoders agnostic to contrast. In both V1 and PPC, stimulus coding rose quickly to high levels within the stimulus period. This information was weaker at low contrasts compared to high contrasts. Choice coding peaked during the spout presentation period after stimulus offset, with weak levels of choice information during the stimulus period. Choice coding varied little with contrast, and there was no obvious difference between V1 and PPC.

We also trained decoders to distinguish Miss trials from CR trials to assess stimulus coding, and from Hit trials to assess stimulus coding (**Figure 5.7D**). The results were qualitatively similar to FA trial decoding, although we did not analyze sessions with varying contrast due to the low number of Miss trials at some individual contrasts. Stimulus information again rose to robust levels within the stimulus period in both regions (**Figure 5.7E**). Choice coding was strongest after stimulus offset, though it began to increase within the stimulus period.

Overall, a population decoding approach did not dramatically improve the level of predictive choice coding compared to that of single neuron analyses (Figure 4.4, Figure 4.6), even in PPC which has been shown to exhibit pre-movement choice signals (Hanks et al., 2015; Harvey et al., 2012; Raposo et al., 2014). It is likely that the asymmetry of the go/no-go design is to blame here, as error trials could result as often from impulsivity or lapses of attention as from errors in

perceptual judgment. Nonetheless, the presence of weak but significant choice signals in PPC is consistent with the previously described error trial analysis (see Figure 2.6) in our variable delay task (Goard et al., 2016).

5.5 Discussion

Understanding how information is encoded and readout from neural activity is a primary goal of systems neuroscience. Although individual neuron response properties can be highly variable, heterogeneous, and dynamic, population-level descriptions can often make better sense of this complexity. Using an encoding and decoding framework, our work has made two major contributions to our understanding of PPC. First, by adding inter-neuronal correlations to an encoding model of V1 and PPC responses, we were able to dramatically improve our ability to predict neuronal activity in the context of a visual decision task. Second, by using a cross-temporal decoding approach we were able to identify differences in the information coding dynamics of V1 and PPC.

5.5.1 Importance of network coupling in encoding models of PPC and V1

Generalized linear models have increasingly been used to describe the statistical dependences of neural activity on measured sensory or task variables, not only in early sensory areas (Pillow et al., 2008) but also in more cognitive areas like PPC (Park et al., 2014; Pinto and Dan, 2015). Our task-only model was able to predict responses to some degree, explaining on average 33% of the variance of V1 responses, but only 20% of the variability in PPC. Interpreting differences or failures in model performance is a challenge for all linear encoding models as they suffer from two major caveats. First, the assumption of linear superposition is strong, particularly when neurons in these areas can have mixed selectivity that may combine different task components in a nonlinear manner. For example, it is possible that the poorer performance of PPC models

could be blamed on an inability to capture nonlinear mixtures of stimulus, engagement, and choice signals. Second, it is never entirely clear whether the chosen explanatory variables are the appropriate ones to use in the first place. For example, PPC neurons may appear to be sensitive to stimulus contrast when it is in fact tuned to a more elusive cognitive variable like decision confidence (Gold and Shadlen, 2007). In addition, it is highly likely that there are other unmeasured variables that drive neural activity which would otherwise be treated as noise.

The additional source of variability that we chose to incorporate was the activity of other neurons in the network. As previously shown by many others (Cohen and Maunsell, 2009; Zohary et al., 1994), a population of neurons can express strong inter-neuronal dependencies that result in correlated activity fluctuations from trial to trial, but see (Ecker et al., 2010; Renart et al., 2010). Our results extend these findings by demonstrating that noise correlations are more prominent in PPC than in V1, are reduced by task engagement, and can improve prediction performance when incorporated into linear models of responses. The possibility that higher cognitive regions may in general express larger correlations than early sensory areas is intriguing and deserves further investigation.

One caveat is that these models and measured correlations were based on calcium responses, an indirect measure of spiking activity. While our lab (Rikhye and Sur, 2015) and others (Cossell et al., 2015; Hofer et al., 2011; Lur et al., 2016) have observed relative changes in correlations using calcium imaging that are qualitatively similar to those obtained from electrophysiological recordings, it must be acknowledged that the absolute strength of correlations cannot be directly related across recording modalities. One particular concern is that contamination of somatic calcium signals by surrounding neuropil can artificially increase the measured mean correlation strength between neurons (Harris et al., 2016). Such errors would lead to a systematic overestimation of correlations, but we observed specific relative differences between PPC and V1, and between Engaged and Passive conditions, which cannot be explained by such overestimation. However, we cannot rule out the possibility that overestimated

correlations due to such artifacts could be partially responsible for the improvement in model performance that we observe.

Another concern is that the lower time resolution of our calcium data makes interpretation of the coupling inputs difficult to interpret. Indeed, previous encoding models using spiking data have described the direct effects of spikes on other neurons on much smaller timescales (Pillow et al., 2008; Truccolo et al., 2005). However, in both cases the coupling interactions are merely phenomenological and should not be interpreted as reflecting biophysical or synaptic mechanisms. Instead, the improvement of performance with the addition of coupling filters (whether at single-spike or spike-rate timescales) reflects the dependence of the responses on network activity, whether it comes from specific sources within the network or a shared external input. In fact, models that incorporate a common input term in addition to inter-neuronal coupling often observe reduced pairwise interactions between neurons (Vidne et al., 2012), further arguing that such models should be taken as being merely descriptive in nature. In our case, this suggests that network fluctuations at longer (hundreds of ms) timescales can account for much of the variability in PPC responses, whether such fluctuations arise from internal or external sources.

5.5.2 Dynamic coding in PPC during task performance

Our analysis of information coding in PPC further suggests that stimulus information can be represented in a dynamic manner in the context of task performance. This contrasts both with V1 and with passive trials in which coding is stable (albeit with a much smaller population of active neurons). Such dynamic population responses have been previously observed in the prefrontal cortex and PPC of primates performing complex decision tasks (Crowe et al., 2010; Meyers et al., 2008; Stokes et al., 2013). Dynamic coding has also been observed in PPC of mice, but only in the context of more complex virtual navigation tasks in which sensory inputs are changing constantly throughout task performance (Harvey et al., 2012; Morcos and Harvey, 2016). Our

data and analyses suggest that such dynamic coding could be a general signature of association cortex, even in the context of more simple non-navigational decision tasks. It also provides evidence that a brain region can switch between dynamic and stable representations depending on task engagement. Finally, preliminary analyses contribute additional evidence that PPC may be able to represent one task-relevant variable (choice) with a dynamic code, while simultaneously encoding another (contrast) using a stationary code (Crowe et al., 2010).

However, a few caveats need to be addressed before definitively making the conclusion that PPC exhibits dynamic coding. First, the width of optimal performance for the across time-bin decoders, while narrower for PPC than V1, are still quite large, on the order of 4 seconds. A major contributor to these large widths is the nature of the calcium signal, which was generated using an indicator with particularly slow offset dynamics, on the timescale of seconds (Chen et al., 2013). Future work could apply appropriate deconvolution or spike inference algorithms (Pnevmatikakis et al., 2016; Theis et al., 2016; Vogelstein et al., 2009) to mitigate this problem, but recent benchmarking tests indicate that even the best available algorithms make substantial errors in predicting spiking activity (Theis et al., 2016). We would argue that our ability to recover some features of dynamic coding even in a low time-resolution calcium dataset suggests that this phenomenon is robust, and that deconvolution would only improve the strength of our findings.

A second, more serious concern is how to reconcile dynamic coding in PPC with the more static delay-period responses observed in PPC during our delayed response task (Goard et al., 2016) (see Figures 2.3, 2.5 for delay-period responses). In this task we were able to summarize the temporal response profiles of task-enhanced PPC neurons with just four distinct types (Figure 2.3), with no neurons responding with transients in the middle of the delay period. Indeed, preliminary cross-temporal decoding analyses of these datasets indicate that although coding is more dynamic during the stimulus and response periods, activity patterns are largely stable during the delay period. One possible explanation of this difference comes from comparisons with the

olfactory system. Studies in the locust antennal lobe have observed that activity patterns encoding odor identity are transiently dynamic at odor onset and offset, but relatively stable during the intervening odor presentation time (Mazor and Laurent, 2005). PPC dynamics in our task may therefore differ from the sequential dynamics observed in virtual navigation tasks (Harvey et al., 2012; Morcos and Harvey, 2016) because changes in sensory input are limited to the stimulus and response periods of our task, in analogy to odor onset and offset in the olfactory system (Mazor and Laurent, 2005). Because the analyses we present in Figure 5.5 come from a 0-s delay task, dynamic coding is expressed throughout the trial, with no intervening stability during the delay period.

PPC thus expresses dynamic patterns of activity during task performance, but does this necessitate mechanisms to readout information in a time-varying manner? Recent work in primates has demonstrated that dynamic coding can coexist with stable readout if there exists a low-dimensional subspace from which information can be decoded (Druckmann and Chklovskii, 2012; Murray et al., 2017). To test for this possibility in our data, we trained decoders to discriminate population activity at multiple time points within a trial, and found that such decoders perform nearly as well across all time points on a held-out test dataset. This suggests that the narrow performance width of the cross-temporal decoders may be due to “overfitting” to the PPC activity patterns in a particular time-bin. Such overfitting did not occur, however, for V1, which had broad performance for decoders trained at most time-bins. This indicates that PPC in our task does have time-varying activity, but it may lie in an intermediate regime between stable activity patterns (as in V1) and truly dynamic codes which would necessitate time-varying readout, as in the case of liquid state machines (Buonomano and Maass, 2009).

5.5.3 Conclusions and outlook

Our results argue that progress in understanding how higher cognitive areas like PPC encode and process information requires a population-level approach. PPC responses are not only dominated by neuronal correlations, but they are also highly dynamic in task-specific ways. These findings were made on the basis of largely separate encoding and decoding analyses. However, encoding models are most successful when they can be leveraged to develop model-based decoders (Park et al., 2014; Pillow et al., 2008), rather than the empirical decoders used here. Before Bayesian model-based decoders can be used, one must demonstrate that the encoding models are successful in predicting responses with a high degree of accuracy. Although such performance levels have been achieved for spiking data using GLMs, accurate modelling of calcium responses remains an active area of research (Lutcke et al., 2013; Mishchenko et al., 2011). Further development of the approaches put forth here for modeling population calcium data holds great promise for yielding deeper insight to how large neuronal populations encode and transform information during behavior.

5.6 Contributions and Acknowledgements

G.N.P. and M.S. conceived the study; G.N.P. performed the experiments; G.N.P. designed and performed the analysis with input from M.S.; G.N.P. wrote the manuscript with input from M.S.

We thank W. Malik, M. Hu, R. Huda, S. Kira, C.D. Harvey, Z. Chen, M.A. Wilson, and M.J. Goard for discussions about the models and analyses used in this chapter. This work was supported by the NIH (M.S., U01-NS090473) and the NSF (G.N.P., Graduate Research Fellowship).

5.7 References

Andersen, R.A., and Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits. Neuron *63*, 568-583.

Averbeck, B.B., Latham, P.E., and Pouget, A. (2006). Neural correlations, population coding and computation. Nat Rev Neurosci *7*, 358-366.

Averbeck, B.B., and Lee, D. (2006). Effects of noise correlations on information encoding and decoding. J Neurophysiol *95*, 3633-3644.

Barbieri, R., Wilson, M.A., Frank, L.M., and Brown, E.N. (2005). An analysis of hippocampal spatio-temporal representations using a Bayesian algorithm for neural spike train decoding. IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society *13*, 131-136.

Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. J Neurosci *31*, 913-921.

Bisley, J.W., and Goldberg, M.E. (2010). Attention, intention, and priority in the parietal lobe. Annual review of neuroscience *33*, 1-21.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Visual neuroscience *13*, 87-100.

Buonomano, D.V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. Nat Rev Neurosci *10*, 113-125.

Chen, J.L., Voigt, F.F., Javadzadeh, M., Krueppel, R., and Helmchen, F. (2016). Long-range population dynamics of anatomically defined neocortical networks. Elife *5*.

Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V.*, et al.* (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295-300.

Cohen, M.R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. Nat Neurosci *14*, 811-819.

Cohen, M.R., and Maunsell, J.H. (2009). Attention improves performance primarily by reducing interneuronal correlations. Nat Neurosci *12*, 1594-1600.

Cortes, C., and Vapnik, V. (1995). Support-Vector Networks. Mach Learn *20*, 273-297.

Cossell, L., Iacaruso, M.F., Muir, D.R., Houlton, R., Sader, E.N., Ko, H., Hofer, S.B., and Mrsic-Flogel, T.D. (2015). Functional organization of excitatory synaptic strength in primary visual cortex. Nature *518*, 399-403.

Crowe, D.A., Averbeck, B.B., and Chafee, M.V. (2010). Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. J Neurosci *30*, 11640-11653.

Druckmann, S., and Chklovskii, D.B. (2012). Neuronal circuits underlying persistent representations despite time varying activity. Curr Biol *22*, 2095-2103.

Ecker, A.S., Berens, P., Keliris, G.A., Bethge, M., Logothetis, N.K., and Tolias, A.S. (2010). Decorrelated neuronal firing in cortical microcircuits. Science *327*, 584-587.

Ecker, A.S., Berens, P., Tolias, A.S., and Bethge, M. (2011). The effect of noise correlations in populations of diversely tuned neurons. J Neurosci *31*, 14272-14283.

Ganguli, S., Bisley, J.W., Roitman, J.D., Shadlen, M.N., Goldberg, M.E., and Miller, K.D. (2008). One-dimensional dynamics of attention and decision making in LIP. Neuron *58*, 15-25.

Goard, M.J., Pho, G.N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife *5*, 558.523.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annual review of neuroscience *30*, 535-574.

Graf, A.B., Kohn, A., Jazayeri, M., and Movshon, J.A. (2011). Decoding the activity of neuronal populations in macaque primary visual cortex. Nat Neurosci *14*, 239-245.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220-223.

Harris, K.D., Quiroga, R.Q., Freeman, J., and Smith, S.L. (2016). Improving data quality in neuronal population recordings. Nat Neurosci *19*, 1165-1174.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62-68.

Hofer, S.B., Ko, H., Pichler, B., Vogelstein, J., Ros, H., Zeng, H., Lein, E., Lesica, N.A., and Mrsic-Flogel, T.D. (2011). Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. Nat Neurosci *14*, 1045-1052.

Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. The Journal of physiology *160*, 106-154.

Huth, A.G., Nishimoto, S., Vu, A.T., and Gallant, J.L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. Neuron *76*, 1210-1224.

Jun, J.K., Miller, P., Hernandez, A., Zainos, A., Lemus, L., Brody, C.D., and Romo, R. (2010). Heterogenous population coding of a short-term memory and decision task. J Neurosci *30*, 916-929.

Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C.E., Kepecs, A., Mainen, Z.F., Qi, X.L., Romo, R., Uchida, N., and Machens, C.K. (2016). Demixed principal component analysis of neural population data. Elife *5*.

Kohn, A., Coen-Cagli, R., Kanitscheider, I., and Pouget, A. (2016). Correlations and Neuronal Population Information. Annu Rev Neurosci *39*, 237-256.

Lur, G., Vinck, M.A., Tang, L., Cardin, J.A., and Higley, M.J. (2016). Projection-Specific Visual Feature Encoding by Layer 5 Cortical Subnetworks. Cell Rep *14*, 2538-2545.

Lutcke, H., Gerhard, F., Zenke, F., Gerstner, W., and Helmchen, F. (2013). Inference of neuronal network spike dynamics and topology from calcium imaging data. Front Neural Circuits *7*, 201.

Machens, C.K., Romo, R., and Brody, C.D. (2010). Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. J Neurosci *30*, 350-360.

Malik, W.Q., Schummers, J., Sur, M., and Brown, E.N. (2011). Denoising two-photon calcium imaging data. PloS one *6*, e20490-e20490.

Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78-84.

Mazor, O., and Laurent, G. (2005). Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. Neuron *48*, 661-673.

Meister, M.L., Hennig, J.A., and Huk, A.C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. J Neurosci *33*, 2254-2267.

Meyers, E.M., Freedman, D.J., Kreiman, G., Miller, E.K., and Poggio, T. (2008). Dynamic population coding of category information in inferior temporal and prefrontal cortex. J Neurophysiol *100*, 1407-1419.

Miri, A., Daie, K., Burdine, R.D., Aksay, E., and Tank, D.W. (2011). Regression-based identification of behavior-encoding neurons during large-scale optical imaging of neural activity at cellular resolution. J Neurophysiol *105*, 964-980.

Mishchenko, Y., Vogelstein, J.T., and Paninski, L. (2011). A Bayesian Approach for Inferring Neuronal Connectivity from Calcium Fluorescent Imaging Data. Ann Appl Stat *5*, 1229-1261.

Morcos, A.S., and Harvey, C.D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. Nat Neurosci *19*, 1672-1681.

Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., and Pouget, A. (2014). Information-limiting correlations. Nat Neurosci *17*, 1410-1417.

Murakami, M., and Mainen, Z.F. (2015). Preparing and selecting actions with neural populations: toward cortical circuit mechanisms. Current opinion in neurobiology *33C*, 40-46.

Murray, J.D., Bernacchia, A., Roy, N.A., Constantinidis, C., Romo, R., and Wang, X.J. (2017). Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. Proceedings of the National Academy of Sciences of the United States of America *114*, 394-399.

Nienborg, H., and Cumming, B.G. (2006). Macaque V2 neurons, but not V1 neurons, show choice-related activity. J Neurosci *26*, 9567-9578.

Paninski, L., Shoham, S., Fellows, M.R., Hatsopoulos, N.G., and Donoghue, J.P. (2004). Superlinear population encoding of dynamic hand trajectory in primary motor cortex. J Neurosci *24*, 8551-8561.

Park, I.M., Meister, M.L., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. Nat Neurosci *17*, 1395-1403.

Pillow, J.W., Paninski, L., Uzzell, V.J., Simoncelli, E.P., and Chichilnisky, E.J. (2005). Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. J Neurosci *25*, 11003-11013.

Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995-999.

Pinto, L., and Dan, Y. (2015). Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed Behavior. Neuron *87*, 437-450.

Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W.*, et al.* (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. Neuron *89*, 285-299.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat Neurosci *17*, 1784-1792.

Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., and Harris, K.D. (2010). The asynchronous state in cortical circuits. Science *327*, 587-590.

Rigotti, M., Barak, O., Warden, M.R., Wang, X.J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature *497*, 585-590.

Rikhye, R.V., and Sur, M. (2015). Spatial Correlations in Natural Scenes Modulate Response Reliability in Mouse Visual Cortex. J Neurosci *35*, 14661-14680.

Rishel, C.A., Huang, G., and Freedman, D.J. (2013). Independent category and spatial encoding in parietal cortex. Neuron *77*, 969-979.

Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci *22*, 9475-9489.

Rorie, A.E., Gao, J., McClelland, J.L., and Newsome, W.T. (2010). Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. PloS one *5*, e9308.

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

Theis, L., Berens, P., Froudarakis, E., Reimer, J., Roman Roson, M., Baden, T., Euler, T., Tolias, A.S., and Bethge, M. (2016). Benchmarking Spike Rate Inference in Population Calcium Imaging. Neuron *90*, 471-482.

Truccolo, W., Eden, U.T., Fellows, M.R., Donoghue, J.P., and Brown, E.N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. J Neurophysiol *93*, 1074-1089.

Vidne, M., Ahmadian, Y., Shlens, J., Pillow, J.W., Kulkarni, J., Litke, A.M., Chichilnisky, E.J., Simoncelli, E., and Paninski, L. (2012). Modeling the impact of common noise inputs on the network activity of retinal ganglion cells. Journal of computational neuroscience *33*, 97-121.

Vogelstein, J.T., Watson, B.O., Packer, A.M., Yuste, R., Jedynak, B., and Paninski, L. (2009). Spike inference from calcium imaging using sequential Monte Carlo methods. Biophys J *97*, 636-655.

Zohary, E., Shadlen, M.N., and Newsome, W.T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. Nature *370*, 140-143.

5.8 Figures

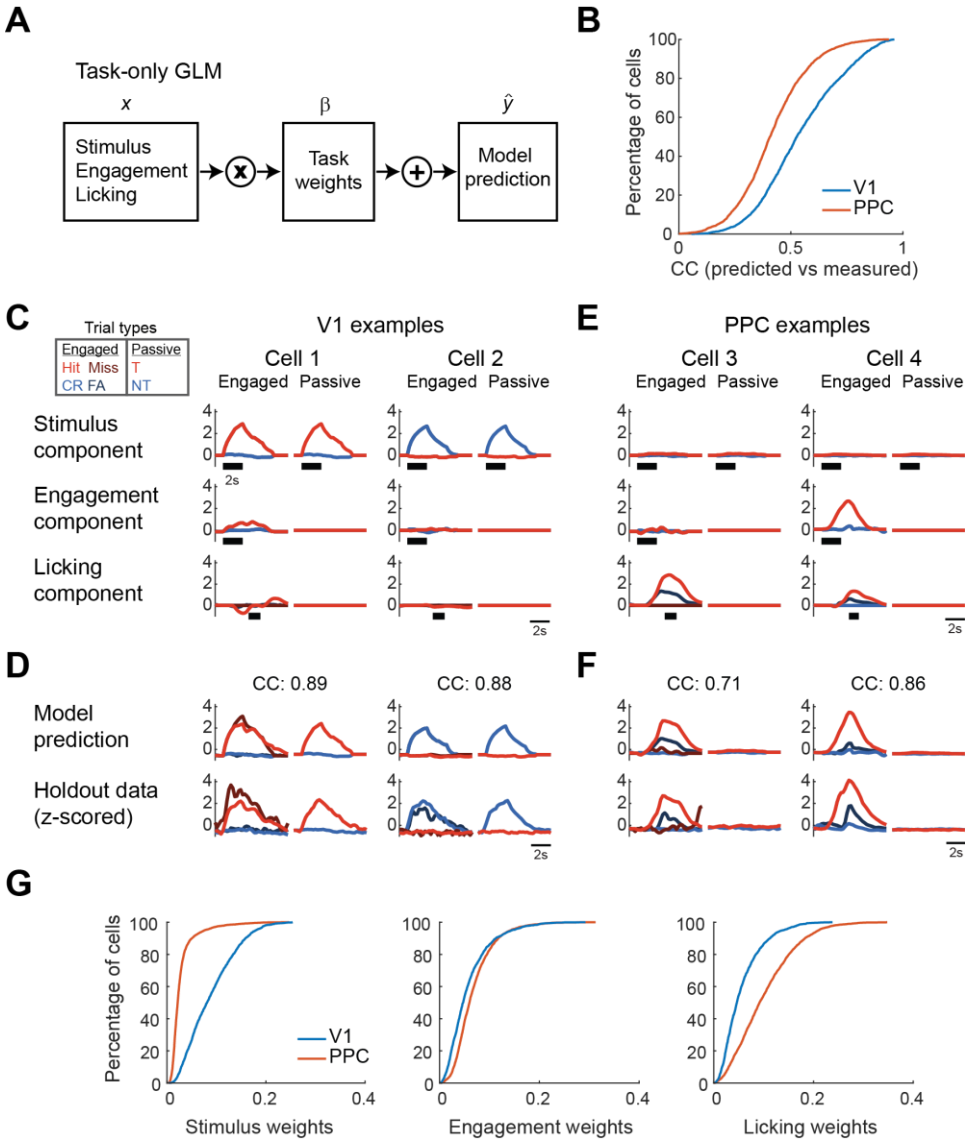

Figure 5.1. Relating V1 and PPC responses to task events using a generalized linear model (GLM).

(**A**) Schematic illustrating GLM formulation for the task-only model. Individual z-scored calcium responses (y) were modeled using lagged predictors for stimulus, engagement, and licking (x). A set of weights (β) were fit for each neuron to generate predictions ($\hat{y}$) on holdout trials. (**B**) Cumulative histogram of prediction performance on holdout trials, quantified for each cell as correlation coefficient (CC) of predictions with actual calcium response. (**C**) GLM model components for two example V1 neurons. These three (Stimulus, Engagement, and Licking) components are summed to generate the model prediction in (D). Trial-averaged components are plotted for six different trial types in Engaged (left) and Passive (right) conditions, with Target trials in red, Non-target trials in blue, and error trials in darker shades. Black bar indicates timing of stimulus (in Stimulus and Engagement components) and licking (in Licking component). Cell 1 and 2 are have large stimulus components for either Target or Non-target stimuli. (**D**) Model predictions and trial-averaged holdout data for Cells 1 and 2. Both cells have high CC. (**E**) Same as (C) but for two PPC cells. Cell 3 has a large licking component, whereas Cell 4 has a larger engagement component. (**F**) Same as (D) but for PPC cells 3 and 4. (**G**) Cumulative histograms of median absolute weight values for V1 and PPC cells. Stimulus weights are larger in V1, whereas licking and engagement weights are larger in PPC.

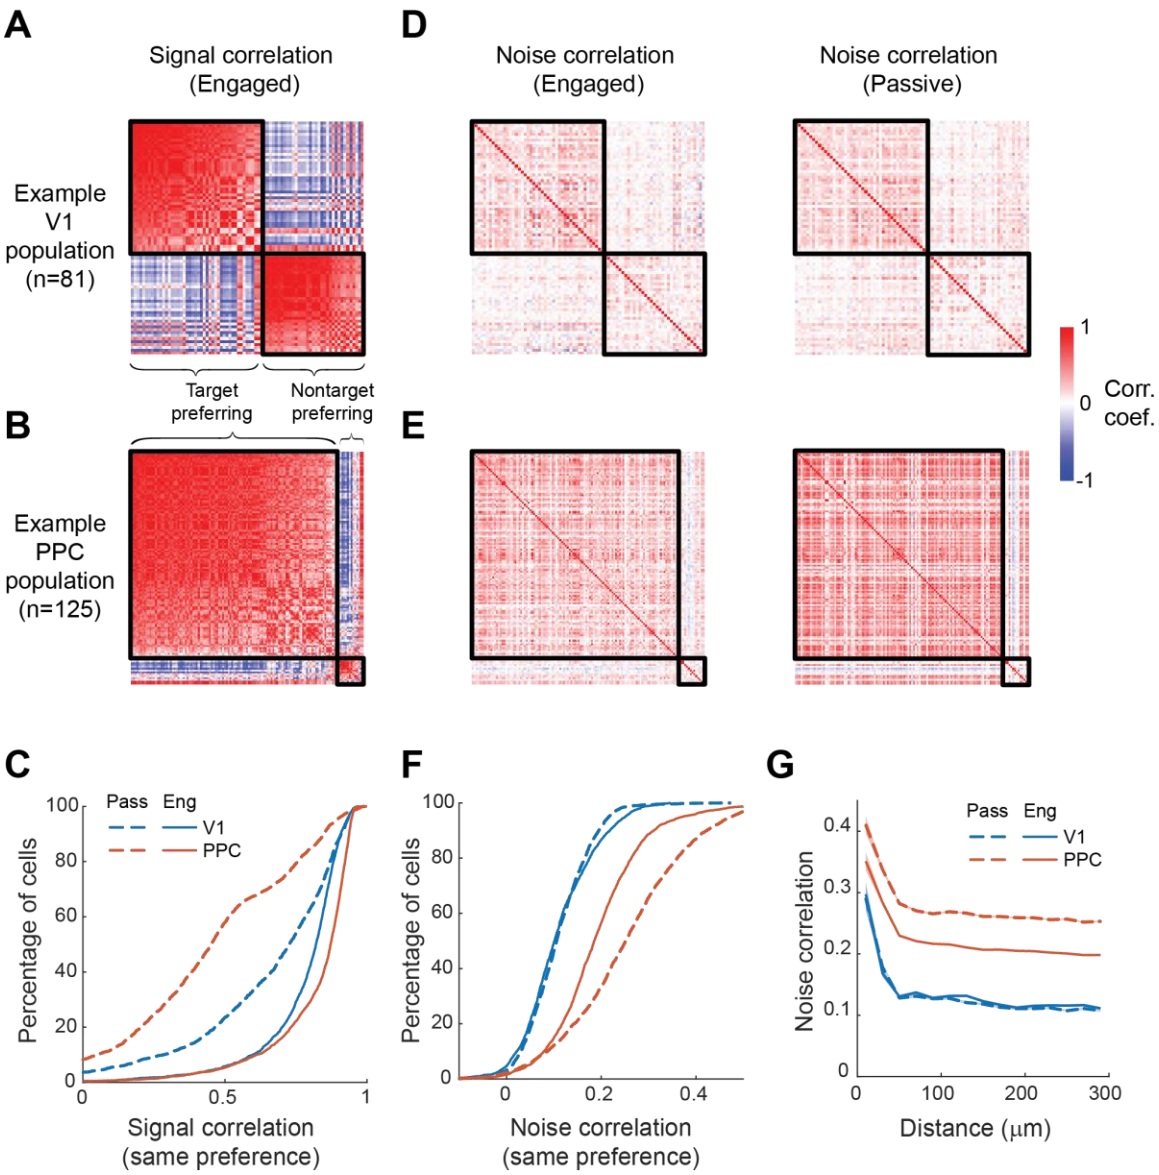

Figure 5.2. Noise correlations are stronger in PPC than in V1.

(**A**) Heatmap showing strength of signal correlations during Engaged trials in a population of V1 cells (n=81). Cells are clustered by stimulus preference, and then sorted by mean within-cluster correlation. (**B**) Same as (A) but for a population of PPC cells (n=125). Most neurons are target-preferring. (**C**) Cumulative histogram of within-cluster signal correlations in V1 (blue) and PPC (red) during Engaged (solid) and Passive (dashed) trials. Signal correlations increase in both areas with engagement, especially in PPC. (**D**) Noise correlations during Engaged (left) and Passive (right) trials in the same population of V1 cells. Within-cluster noise correlations are largely unchanged with engagement. (**E**) Same as (D) but for a population of PPC cells. Within-cluster noise correlations are reduced with engagement. (**F**) Cumulative histogram of within-cluster noise correlations in V1 and PPC during Engaged and Passive trials. Noise correlations are higher in PPC, especially during Passive trials. (**G**) Noise correlations as a function of between-cell distance in V1 and PPC during Engaged and Passive trials. Cells tend to have stronger noise correlation within 50μm, but the distance-dependence is reduced beyond 100μm.

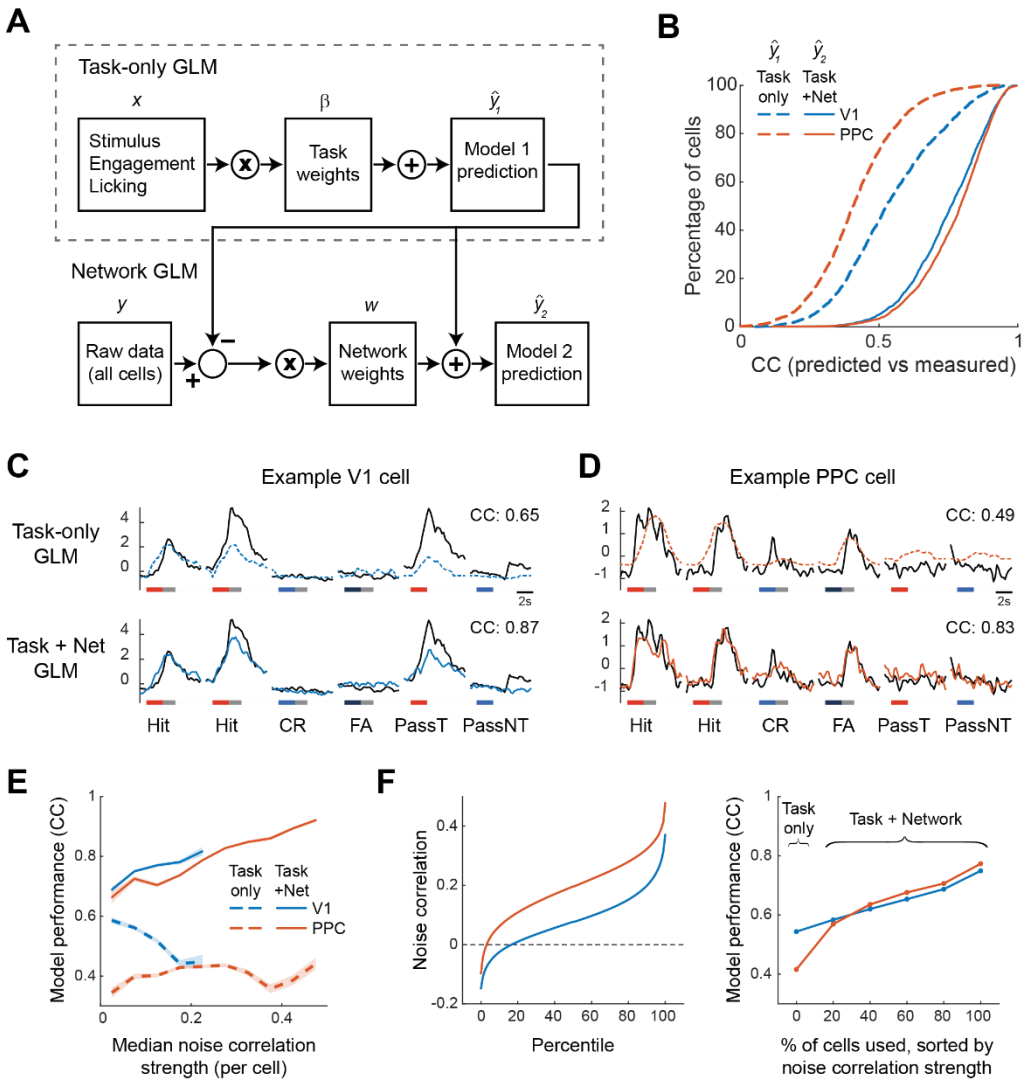

Figure 5.3. Addition of network coupling terms in GLM improves model prediction.

(**A**) Schematic illustrating GLM formulation for the Task + Network GLM (Model 2). The network GLM is fit using residuals from the Task-only GLM (Model 1). Network weights (w) are computed by fitting the residuals of one cell to a weighted sum of residuals from all other simultaneously recorded cells in the network. The predicted residuals for a given cell are then added to the Model 1 prediction to generate a new prediction (Model 2, $\hat{y}_2$). (**B**) Cumulative histogram of prediction performance on holdout trials in V1 (blue) and PPC (red) for both the Task-only model (dashed, same as Figure 5.1B) and the Task + Network model (solid). Prediction performance improves for both areas with Model 2, but especially for PPC. (**C**) Comparison of performance of Task-only model (top, dashed), and Task + Network model (bottom, solid) for one example V1 cell on six holdout trials. Raw z-scored data is in black, and model prediction is colored. Colored bars indicated stimulus timing and type (Target is red, Non-target is blue), and gray bar indicates spout presentation. (**D**) Same as (C) but for one PPC cell. (**E**) Task + Network model performance is correlated with the strength of noise correlations for a given cell. Cells were binned based on their median absolute noise correlations with other cells, and shading indicates SEM. For V1, task-only model performance was inversely correlated with noise correlation strength. (**F**) Network inputs were sorted by noise correlation (left) in ascending order to train models on subsets of the network. Performance of models (right) trained using subsets with the weakest noise correlations (20, 40, 60, and 80% of cells) as compared with task-only (0%) and full network (100%) models.

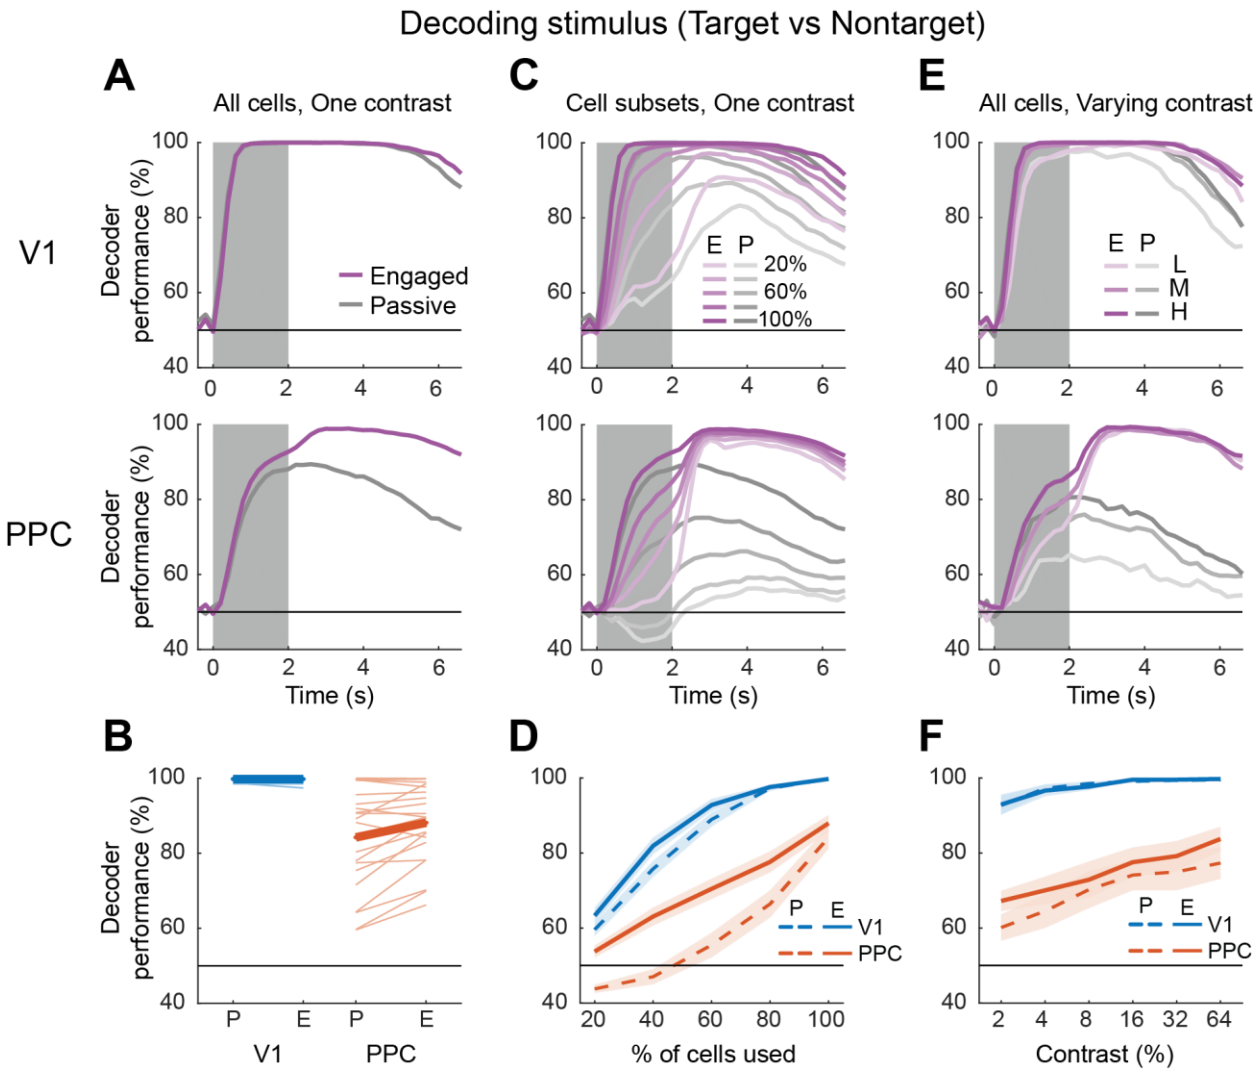

Figure 5.4. Decoding stimulus from V1 and PPC during engaged and passive conditions.

Population decoders (linear SVM) were trained to discriminate target from non-target stimuli using passive trials or correct engaged trials. All plots represent average cross-validated performance across different V1 and PPC populations. Significant ($p<0.05$) decoding accuracy was achieved at 58%. (**A**) Separate decoders for Engaged (purple) and Passive (gray) trials were trained using all cells in each V1 (n=18 populations) and PPC (n=22) population to discriminate Target versus Nontarget trials at each time bin. (**B**) Average decoder performance during the stimulus period (0.8 to 2 s) for each individual V1 and PPC population (thin lines), and on average (thick lines), during Passive (P) and Engaged (E) conditions. (**C**) Same as (A), but decoders were trained on increasing numbers of neurons, added from least to most selective (shade indicates 20, 40, 60, 80 or 100% of neurons), in either Engaged or Passive conditions. Using the 60% least selective PPC neurons on Passive trials leads to chance levels of prediction (bootstrapped chance level is 58% at $p<0.05$). (**D**) Average performance during the stimulus period for different subsets of neurons in V1 (blue) and PPC (red), during Engaged (solid) or Passive (dashed) conditions. (**E**) Same as (A), but decoders were trained on V1 (n=8) and PPC (n=11) datasets in which stimulus contrast was varied from trial-to-trial. Decoders were trained without information about contrast, but performance is plotted as a function of contrast (shade indicates Low, Medium or High contrast). (**F**) Average performance during the stimulus period for different stimulus contrasts in V1 or PPC, during Engaged and Passive conditions.

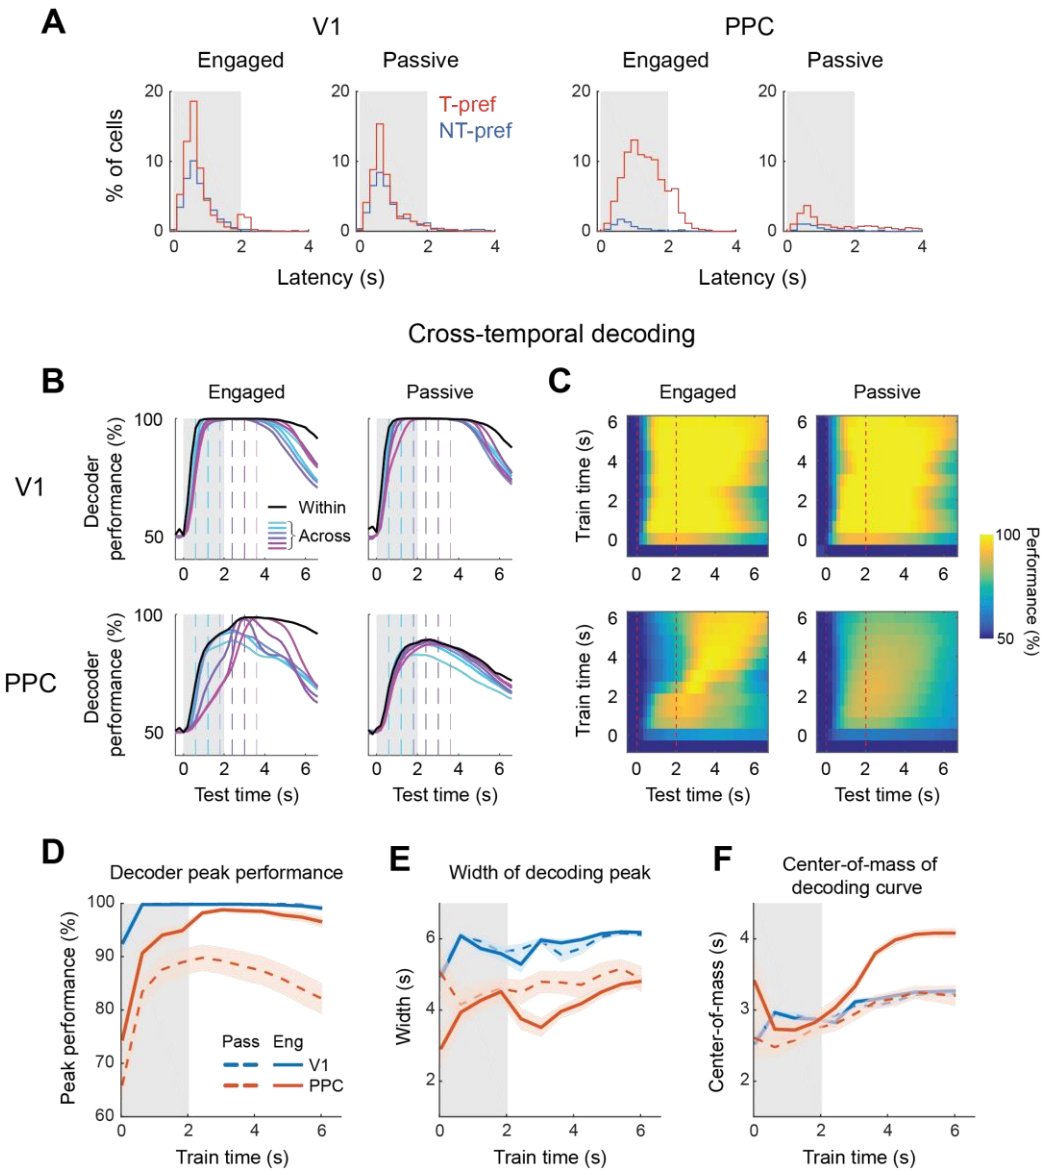

Figure 5.5. Time-dependent decoding in PPC but not V1.

(**A**) Histogram of response onset latency, relative to stimulus onset, in V1 and PPC. Cells are separated into target- (red) and nontarget-preferring (blue), and only cells with significant responses during Engaged or Passive trials are counted. PPC exhibits a much broader distribution of latencies in V1, specifically during Engaged trials. (**B-F**) To test for time-dependent coding, a cross-temporal decoding approach was used. Decoders were trained to distinguish population activity during Target versus Nontarget trials at one time bin, and then tested across all time bins. (**B**) Performance of decoders in V1 and PPC during Engaged and Passive trials. "Within" time bin decoding performance is plotted in black (same as Figure 5.4A). "Across" time bin decoding performance is plotted in color, with vertical dashed lines indicated training time of each decoder. (**C**) Heatmaps of decoder performance, with each row representing a decoder trained at one time bin. Red dashed lines indicate stimulus onset and offset. (**D**) Peak performance of decoders in V1 (blue) and PPC (red) during Engaged (solid) and Passive (dashed) trials, as a function of training bin. (**E**) Width of decoding peak (above a performance threshold of 75%), as a function of training bin. Width is consistently narrower for PPC compared to V1. (**F**) Center-of-mass of decoding peak as a function of training bin. Center-of-mass does not change much across training bins, except for PPC during Engaged trials.

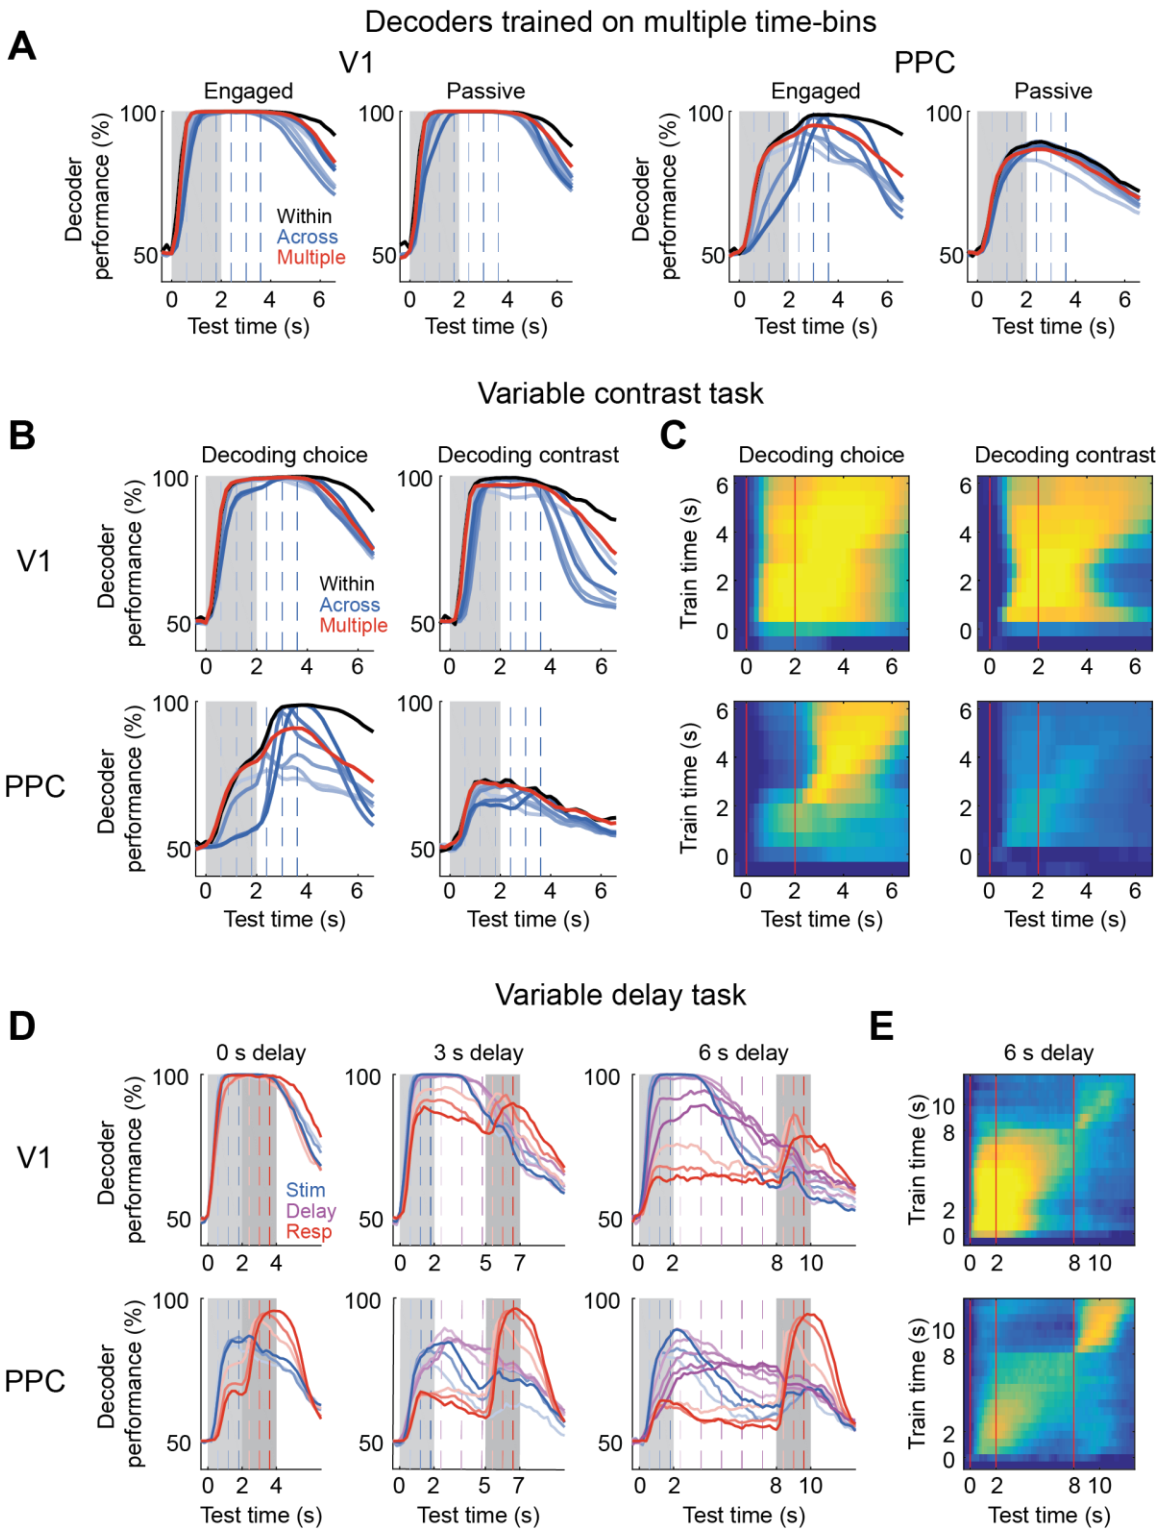

Figure 5.6. Time-dependent coding during variable contrast and variable delay tasks.

(**A**) A decoder was trained to discriminate Hit versus Correct Reject trials on multiple time-bins (red) for both V1 and PPC during Engaged and Passive conditions. The “Multiple” time-bin decoder was trained using the population vectors across six time-bins (blue dashed lines), and then tested on all time-bins of a

validation set. “Within” (black) and “Across” (blue) time-bin decoders are identical to those presented in Figure 5.5B. Blue dashed lines also indicate training time for individual “Across” time-bin decoders. (**B-C**) Decoding choice and stimulus contrast in the variable contrast task. Decoders for choice discriminated Hit and Correct Reject trials and ignored contrast information. Decoders for contrast discriminated High contrast Hit trials from Low Contrast Hit trials (**B**) Performance of “Within” (black), “Across”, and “Multiple” decoders in V1 and PPC. Decoding performance appears dynamic in PPC for choice but not for contrast. (**C**) Heatmap of decoder performance for both choice and contrast in V1 and PPC, with each row representing a decoder trained at a particular time-bin. Red lines indicate stimulus onset and offset. (**D-E**) Decoding choice (Hit versus Correct Reject) in the variable delay task. Separate decoders were trained for each delay duration. (**D**) Performance of “Across” time-bin decoders trained during the stimulus period (blue), during the delay period (purple), and during the response period (red), for 0 s, 3 s, and 6 s delay trials. Dashed lines indicate decoder training time. Gray shading indicates the stimulus period and the response period. (**E**) Heatmap of decoder performance for 6 s delay trials. Red lines indicate stimulus onset, stimulus offset, and spout onset (beginning of response period).

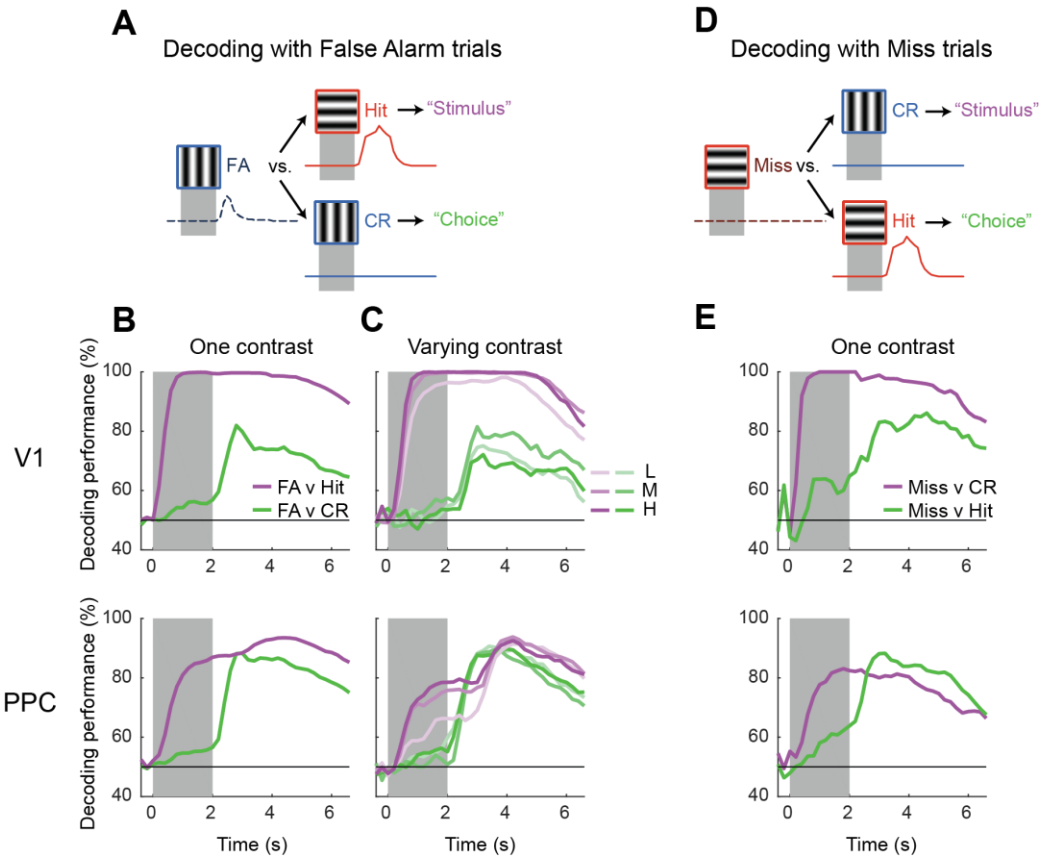

Figure 5.7. Decoding stimulus and choice using error trials.

All population decoders in this figure were trained and tested at the same time bins. (**A**) Decoders were trained to discriminate FA trials from Hit trials to assess "Stimulus" encoding (purple), or to discriminate FA trials with CR trials to assess "Choice" encoding (green). (**B-C**) Performance of stimulus and choice decoders using FA trials in V1 and PPC. Significant ($p<0.05$) decoding accuracy was achieved at ~66%. (**B**) Decoders were trained using sessions in which one contrast was used. (**C**) Decoders were trained using sessions in which multiple contrasts were used. Decoders were trained without information about contrast, but performance is plotted as a function of contrast (shade indicates Low, Medium or High contrast). (**D**) Decoders were trained to discriminate Miss trials from CR trials to assess "Stimulus" encoding (purple), or to discriminate Miss trials with Hit trials to assess "Choice" encoding (green). (**E**) Same as (**B**) but using Miss trials instead of FA trials. Significant ($p<0.05$) decoding accuracy was achieved at ~74%. No equivalent analysis to (C) was performed for Miss trials due to low numbers of trials at individual contrasts.

Chapter 6: Discussion

6.1 Summary of key findings

In this thesis we investigated the neural mechanisms underlying a visual decision task in mice by recording and manipulating neural activity in three cortical areas: primary visual cortex (V1), posterior parietal cortex (PPC), and frontal motor cortex (fMC). We asked three major questions in the context of this task: 1) How is information relevant to a decision encoded across the cortex? 2) What cortical areas are necessary for perceptual decision-making? And finally, 3) What neural mechanisms underlie the process of mapping sensory percepts to appropriate motor outputs?

6.1.1 How is information relevant to a decision encoded across the cortex?

Experiments in **Chapter 2** yielded two key insights towards answering this question. First, neurons in sensory, association, and motor cortices can have widely heterogeneous response properties. While the majority of V1 neurons respond during stimulus presentation and the majority of fMC neurons respond during motor output, there was a significant proportion of neurons in each area that responded during the other task epochs. PPC neurons were the most heterogeneous, as explored in further detail in **Chapter 4**, containing both contrast-dependent sensory and contrast-insensitive motor neurons.

Second, despite this heterogeneity, each region had distinct encoding dynamics. V1 primarily encoded the stimulus, and time-dependent decoding analyses in **Chapter 5** revealed that this coding is stable with time. PPC encoded the stimulus at first, which was rapidly transformed into information about choice. Time-dependent decoding revealed that PPC can exhibit dynamic coding, but whether this information primarily represents stimulus or choice is unclear. However, this dynamic code was specific to performance of the task, suggesting that it

may be more related to the animal's choice. Finally, information in fMC primarily encoded choice, even though it arose quite early during the stimulus period. The encoding analyses suggest that sensorimotor transformation can occur rapidly in this delayed response task.

While this work establishes the encoding properties of these three areas in a delayed response task, future work remains to determine whether similar signals are present in other visuomotor areas of the mouse. Other areas of interest include secondary visual areas (Makino and Komiyama, 2015), the prefrontal cortex (Zhang et al., 2014), as well as subcortical regions like the striatum (Znamenskiy and Zador, 2013) and superior colliculus (Duan et al., 2015). Additionally, future experiments should be performed using more sophisticated decision tasks (e.g. forced-choice instead of go/no-go) that allow for better isolation of signals relating to stimulus, choice, attention and reward.

6.1.2 What cortical areas are necessary for perceptual decision-making?

Optogenetic experiments in **Chapter 3** established the necessity of each area in performance of the task. All three areas were necessary for performance of the task during the stimulus period, even with very short (250 ms) stimulus presentations. The necessity of V1 is consistent with its known role in processing orientation. However, the role of PPC is difficult to assess independently from V1, as we demonstrated using simultaneous recording and inactivation. Whether PPC plays an instructive versus merely permissive role could be definitely established in future work using pharmacological lesion experiments. The early involvement of fMC is consistent with a rapid sensorimotor transformation process. However, it is also possible that similar to PPC, fMC provides permissive signals to other areas that are critical for the sensorimotor transformation.

The frontal motor cortex was the only area tested that was necessary during the delay and response epochs. Whether the memorized motor plan is maintained solely within the fMC, or

whether it participates with other regions during the delay period is unclear. Additionally, the involvement of other areas like PPC may be necessary when stimulus information must be kept in memory, as in the case of a delayed match-to-sample task. Nonetheless, our results point to a critical role for the frontal motor cortex is maintaining and executing decisions in a delayed response task (Liu et al., 2014). Future work remains to identify whether other cortical (Erlich et al., 2015) or subcortical (Duan et al., 2015) areas are causally involved in this and similar decision tasks.

6.1.3 What neural mechanisms underlie the process of mapping sensory percepts to appropriate motor outputs?

To address the question of sensorimotor transformation, we focused on the posterior parietal cortex using a combination of behavioral manipulations and computational tools (**Chapter 4** and **Chapter 5**). The strong task dependence of PPC responses argue against a primary role in sensory processing, but do not prove its role in sensorimotor transformation. Additional evidence comes from the presence of heterogeneous sensory, motor, and sensorimotor neurons, intermingled within PPC. One speculative hypothesis is that this heterogeneity reflects a sensorimotor transformation occurring between separate subpopulations of PPC neurons. However, a more parsimonious explanation may be that PPC, as shown in other higher cortical regions, encodes multiple task-relevant signals simultaneously that can be read-out separately as needed by downstream circuits.

Our strongest evidence for PPC's role in sensorimotor transformations comes from the reversal experiments, in which most PPC neurons reflect the learned sensorimotor contingency. However, this experiment does not show that PPC neurons are actually involved in the sensorimotor re-mapping process. Future work remains to demonstrate whether PPC and/or

other areas including prefrontal cortex or the basal ganglia (Pasupathy and Miller, 2005) are critical for the sensorimotor transformation, and what mechanisms underlie the process itself.

6.2 Future directions

It is remarkable what can be learned using the modern tools of population imaging and optogenetic manipulation in the context of even a very simple visual decision task. Experiments like ours, which involve measurement and perturbation of multiple brain regions, were once prohibitively difficult, but are now accessible to many laboratories interested in the neurobiology of perceptual decisions. This is what the mouse system and our modern age of neurotechnology promises (O'Connor et al., 2009): recordings from more neurons across more areas (Sofroniew et al., 2016), distinctions between molecularly identified cell-types (Pinto and Dan, 2015), projection-specific recording and manipulation (Chen et al., 2013; Wimmer et al., 2015), and even targeted perturbation of functionally identified ensembles (Liu et al., 2012).

However, as exciting as these novel tools are, I would argue that deeper insights into the mechanisms of perceptual decisions will not achieved primarily by utilizing more and more powerful and specific tools. Instead, achieving true understanding into how the brain performs a behavior must start with a deep understanding (Krakauer et al., 2017) and experimental control (Fetsch, 2016) of the behavior itself. While many of the limitations of this thesis come from the simplicity of the task design, the future is bright for combining our multi-areal interrogation approach with rich and carefully designed experimental tasks (Brunton et al., 2013), and higher dimensional descriptions of the behavior itself (Wiltschko et al., 2015).

Finally, as our capacity to record from more and more neurons increases, we must develop stronger theoretical and computational frameworks for interpreting such data (Cunningham and Yu, 2014; Gao and Ganguli, 2015). Developing a conceptual understanding of how the brain

performs behavior cannot be achieved by merely gathering more and better data. After all, even for networks in which we do have near complete information, such as artificial neural networks for deep learning (Gao and Ganguli, 2015) or even the humble microprocessor (Jonas and Kording, 2017), meaningful understanding of their function is not achievable with current bottom-up analytical approaches. Better theoretical approaches are therefore more important than ever, and we in particular need frameworks that can provide specific predictions to guide experimental intervention (Panzeri et al., 2017). By extending the approach of this thesis to combine richer behavioral descriptions, computational theory, and interventional experiments, we as a field can hope to arrive at a deeper conceptual understanding of how the brain controls behavior.

6.3 References

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95-98.

Chen, J.L., Carta, S., Soldado-Magraner, J., Schneider, B.L., and Helmchen, F. (2013). Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. Nature *499*, 336-340.

Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. Nat Neurosci *17*, 1500-1509.

Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Requirement of Prefrontal and Midbrain Regions for Rapid Executive Control of Behavior in the Rat. Neuron *86*, 1491-1503.

Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. Elife *4*.

Fetsch, C.R. (2016). The importance of task design and behavioral control for understanding the neural basis of cognitive functions. Current opinion in neurobiology *37*, 16-22.

Gao, P., and Ganguli, S. (2015). On simplicity and complexity in the brave new world of large-scale neuroscience. Current opinion in neurobiology *32*, 148-155.

Jonas, E., and Kording, K.P. (2017). Could a Neuroscientist Understand a Microprocessor? PLoS Comput Biol *13*, e1005268.

Krakauer, J.W., Ghazanfar, A.A., Gomez-Marin, A., MacIver, M.A., and Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. Neuron *93*, 480-490.

Liu, D., Gu, X., Zhu, J., Zhang, X., Han, Z., Yan, W., Cheng, Q., Hao, J., Fan, H., Hou, R.*, et al.* (2014). Medial prefrontal activity during delay period contributes to learning of a working memory task. Science *346*, 458-463.

Liu, X., Ramirez, S., Pang, P.T., Puryear, C.B., Govindarajan, A., Deisseroth, K., and Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. Nature *484*, 381-385.

Makino, H., and Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in the visual cortex. Nat Neurosci *18*, 1116-1122.

O'Connor, D.H., Huber, D., and Svoboda, K. (2009). Reverse engineering the mouse brain. Nature *461*, 923-929.

Panzeri, S., Harvey, C.D., Piasini, E., Latham, P.E., and Fellin, T. (2017). Cracking the Neural Code for Sensory Perception by Combining Statistics, Intervention, and Behavior. Neuron *93*, 491-507.

Pasupathy, A., and Miller, E.K. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. Nature *433*, 873-876.

Pinto, L., and Dan, Y. (2015). Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed Behavior. Neuron *87*, 437-450.

Sofroniew, N.J., Flickinger, D., King, J., and Svoboda, K. (2016). A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. Elife *5*.

Wiltschko, A.B., Johnson, M.J., Iurilli, G., Peterson, R.E., Katon, J.M., Pashkovski, S.L., Abraira, V.E., Adams, R.P., and Datta, S.R. (2015). Mapping Sub-Second Structure in Mouse Behavior. Neuron *88*, 1121-1135.

Wimmer, R.D., Schmitt, L.I., Davidson, T.J., Nakajima, M., Deisseroth, K., and Halassa, M.M. (2015). Thalamic control of sensory selection in divided attention. Nature *526*, 705-709.

Zhang, S., Xu, M., Kamigaki, T., Hoang Do, J.P., Chang, W.C., Jenvay, S., Miyamichi, K., Luo, L., and Dan, Y. (2014). Selective attention. Long-range and local circuits for top-down modulation of visual cortex processing. Science *345*, 660-665.

Znamenskiy, P., and Zador, A.M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. Nature *497*, 482-485.

Supplementary Materials

Behavioral training stages

Mice were trained once a day, 5-6 days per week. Mice were trained in successive stages, with advancement to the next stage contingent on correct performance (**Table 1**): 1) Mice receiving reward any time they licked the spout. 2) Trial structure was initiated by having an auditory cue tone, followed by a visual stimulus (100% targets), followed by an inter-trial interval. Mice were only rewarded for licks during the visual stimulus. 3) Once mice exhibited preferential licking during the stimulus, the target rate was reduced over several sessions from 100% to 50%. At this point, the non-target was a static grating orientated orthogonally to the target. Licks during non-targets were punished with white noise or white noise plus quinine. 4) Once mice exhibited the ability to discriminate target and non-target gratings ($d' > 1$ and $R_{HIT} - R_{FA} > 30\%$ for several sessions, where R_{HIT} and R_{FA} are the hit and false alarm rate, respectively), the temporal frequency of the non-target grating was increased. 5) Spout withdrawal was introduced. At first the spout was extended within range before the stimulus appeared, then spout extend time was gradually delayed until after the stimulus had turned off. 6) Some mice were also trained on the variable contrast version of the task. Stimuli of lower contrast were gradually added and randomly interleaved in with higher contrast stimuli, until mice could achieve significant performance even at the lowest contrast (2%). Mice that failed to fully learn the task within 150 sessions or showed signs of infection were removed from the study. Some mice were additionally trained on a delayed response version of the task which added to the training time.

Video analysis of movement during delay period

We performed a video analysis (**Figure S1**) to show that mice did not exhibit postural changes or increased movement during the delay period (Goard et al., 2016). Cropped video frames (300 x

200 pixels; width x height) from Hit and CR trials were compared to a “template” CR image to measure postural changes or changes (increases or decreases) in movement during each epoch of the task (Pre-stimulus, Delay, Response). Since some amount of movement is expected in all conditions, a pixel-wise map of the absolute difference between single CR frames and the CR template within-condition (D_{CR}) was calculated as a measure of baseline movement:

$$D_{CR}(x,y) = \left|CR_f(x,y) - \overline{CR_{F\neq f}(x,y)}\right|$$

Where *f* is the index of a single CR frame and ***F≠f*** is the set of all CR frames except *f*, and where *x* and *y* are pixel indices. The absolute difference map was calculated separately for each epoch (Pre-stimulus, Delay, Response). A pixel-wise map of the absolute difference between single Hit frames and the CR template (D_{Hit}) was calculated in the same manner:

$$D_{Hit}(x,y) = \left|Hit_f(x,y) - \overline{CR_{F\neq f}(x,y)}\right|$$

In cases where the number of Hit frames exceeded the number of CR frames, the excluded frame was chosen at random from the CR frames. Finally, the difference in movement on Hit trials relative to CR trials (D_{Sub}) was calculated by taking the absolute value of the subtracted difference maps (**Figure S1B**):

$$D_{Sub}(x,y) = \left|D_{Hit}(x,y) - D_{CR}(x,y)\right|$$

Note that since the frames are compared against a CR template, this approach will capture not only transient movement but also stable postural changes specific to Hit trials. To compare between sessions, the subtracted difference maps (D_{Sub}) were averaged across all pixels for each epoch (**Figure S1C**).

References

Goard, M.J., Pho, G.N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife *5*, 558.523.

Supplementary Figures and Tables

Stage	**Goal**	**Sessions (Mean ± SEM)**
Habituation and lick training	Acclimate to licking while head-fixed	3.9 ± 0.3
Target stimuli only	Introduce trial structure; Lick at appropriate time after stimulus	6.6 ± 1.0
Static non-targets (from 25% to 50% of stimuli)	Discriminate target from non-targets	11.7 ± 1.7
Moving non-targets (from 0 Hz to 2 Hz temporal frequency)	Discriminate target from non-targets	28.3 ± 4.0
Spout withdrawal (from 1 second before to 2 seconds after stimulus onset)	Wait until spout extends to lick	21.2 ± 5.0
Variable contrast (from medium contrasts 16-64% to low 2-64%)	Discriminate at low and variable contrasts	4.8 ± 2.6
Contingency reversal (immediate)	Lick to new target and suppress licking to new non-target	8.7 ± 0.3

Table S1

Description of different behavioral training stages, their goals, and average training time (daily sessions, usually 5 or 6 days a week). Not all mice were advanced to the last two stages (variable contrast and contingency reversal).

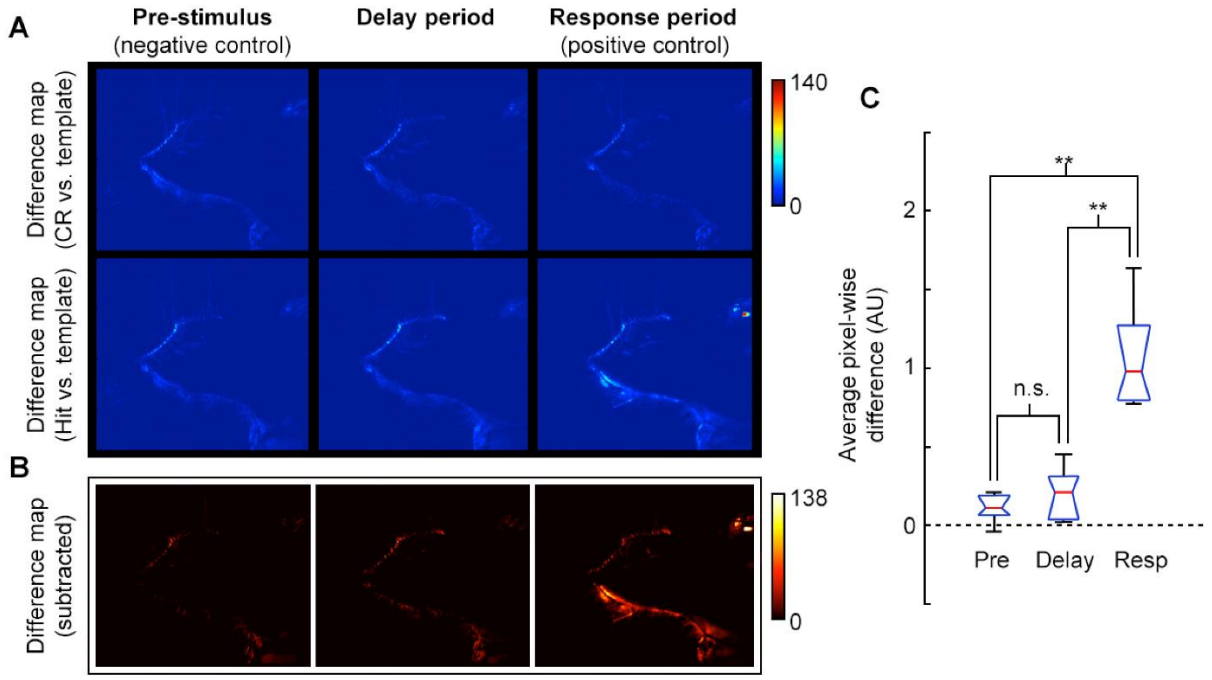

Figure S1: Video analysis of movement during delay period.

(**A**) Pixel-wise absolute difference map between CR template (average of CR frames) and individual CR frames (top row) or Hit frames (bottom) for pre-stimulus, delay, and response epochs of task. Movements relative to template will appear as bright spots on the difference map. (**B**) Absolute difference of CR-template difference map (A, top row) from Hit-template difference map (A, bottom row) to reveal pixels that changed during Hit trials. Movement or postural adjustments during Hit trials will appear as hot spots on the difference map (e.g., see hot spots near mouth, eye, and ear for Response epoch). (**C**) Box plot (red, median; blue, quartile; black, range) of average pixel-wise differences for pre-stimulus epoch (Pre), delay epoch (Delay) and response epoch (Resp) across sessions. The average pixel-wise difference during the Delay epoch was not significantly different from the Pre epoch, ($p = 0.39$), but both Pre and Delay were significantly different from the Response epoch ($p = 0.005$ and $p = 0.009$, respectively; t-test; $n = 5$ sessions, consisting of 3040 frames from 41 trials across 3 mice; all mice at plateau performance levels).