UNIVERSITY OF CALIFORNIA, SAN DIEGO

Perturbation of the Parietal Motor Network in a Sensorimotor Discrimination Task

A Thesis submitted in partial satisfaction of the requirements for the degree
Master of Science

in

Biology

by

Madan Mukundan

Committee in charge:

Professor Takaki Komiyama, Chair
Professor Roberto Malinow, Co-Chair
Professor Nicholas Spitzer

2016

The Thesis of Madan Mukundan is approved and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

Co-Chair

_____

Chair

University of California, San Diego

2016

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

I'm incredibly grateful for the time I've had the opportunity to spend working here in the lab. When I was first joined, I couldn't have possibly had any idea of how much I would change as a person as a result of my time here. I've learned more about work, science, and life than I could have possibly imagined, and I still sometimes have a hard time thinking about how much I've changed.

There are several people that I would like to thank in the lab who have helped me through thick and thin in my time here, starting with:

EunJung Hwang, for bringing me into the lab, trusting me to help her with her work, and guiding and mentoring me in my time here. Her scientific brilliance and capacity for hard work are matched only by her kindness, and working with her has been nothing but a pleasure.

Takaki Komiyama, my advisor and committee chair, for giving me the opportunity to continue working in the lab and pursue the BS/MS program, for pushing me to work harder in the lab, and for teaching me what it is like to work in science. I genuinely wish I sought his guidance more often than I did. I know the lab will find nothing but success in the future.

Jun Lee, a former post-doc in the lab, for teaching me almost all of the lab techniques I know now, guiding my scientific thinking in my time here, and giving me the opportunity to continue working in the future. He was an amazing friend and a mentor to me in a time when I needed it most, and I am eternally grateful to him.

Jeff Dahlen, for helping me with a mind-numbing amount of work, for teaching me to navigate the workings of Matlab (which I still have trouble with), for pulling me through every time I panicked, and for helping me anyway every time I didn't listen to him.

Anna Kim for teaching me sectioning and histology, and for all of the work she has done for all of our projects in the lab.

And of course, the rest of the lab members and former lab members who have made incredibly strong impressions on me, for whom I am very grateful to have gotten to know: Nathan Hedrick, Neehar Kondapaneni, Lovie Li, Monica Chu, Keelin O'Neil, Haixin Lu, Akinori Mitani, Chi Ren, Andrew Peters, Ryoma Hattori, Bethanny Danskin, An Wu, Shan Lu, and all of our undergraduates who assist us in running animals every single day. I am sorry I can't write for you all individually.

I would also like to thank Roberto Malinow, my committee co-chair, and Massimo Scanziani for introducing me to neurobiology as an undergraduate and inspiring me to go into neuroscience research.

I'd like to thank Nick Spitzer and Roberto Malinow for being on my committee, and for helping guiding me through this process.

I would like to thank my parents and my brother for helping and supporting me throughout my whole life, and during my time here.

Last, but arguably most importantly, I would like to thank the hundreds of mice who have given their lives for the progression of my science, and the many more animals who are pushing the field forward today.

Jeff Dahlen created the stereotypy of movement plots in Figure 1.1D and E, Figure 1.2G and the corresponding part of results in Chapter 1.

Chapter 2 (specifically part of Results, Figure 2.1, 2.2, 2.3, and S1, and part of Methods) is adapted from material that has been submitted for publication. Hwang, E.J., Dahlen, J.E., Mukundan, M., Komiyama, T. (2016). History-based Action Selection Bias in Posterior Parietal Cortex. The thesis author was a co-author of this material.

ABSTRACT OF THE THESIS

Perturbation of the Parietal Motor Network in a Sensorimotor Discrimination Task

by

Madan Mukundan

Master of Science in Biology

University of California, San Diego, 2016

Professor Takaki Komiyama, Chair

Professor Roberto Malinow, Co-Chair

The core function of the brain is to make meaningful decisions based on environmental stimuli. While the basic response properties of sensory regions have largely been agreed upon in recent times, a driving force of neuroscience research is how sensory information is integrated with biases, motivations, and assessments of internal state, and then transformed into robust choices and movements.

The primary motor cortex (M1) has long been known as the main controller of movement for some time, but the exact role of M1 and how it communicates with

movement preparation-associated areas, like the secondary motor cortex (M2) and posterior parietal cortex (PPC), are still being elucidated. Together, these three areas form the parietal motor network, and their study is essential to understanding the sensorimotor transformation.

In the first experiment, mice were trained in a novel two-alternative forced-choice visual discrimination task, the requirement of these areas was probed through optogenetic inactivation. M1 inactivation caused animals to default choice to their more stereotyped direction. M2 inactivation caused similar deficits in discrimination, and also impaired hold performance. PPC was found to be non-essential for the task.

PPC is known to be involved in aspects of decision-making, motor coordination, and analysis of value. In the second experiment, mice trained to perform a visually-guided memory task developed a reward-choice history-dependent bias revealed through mathematical modeling. Optogenetic perturbation of PPC activity during the pre-stimulus, but not during the post-stimulus period, altered the internal bias. As a result, the model's predictive capability was significantly lowered in inactivation trials.

Chapter 1: Inactivating the Parietal Motor Network in a Visual Discrimination Task

Introduction

The Sensorimotor Transformation

The ability to act on the world is essential to survival for all forms of life on Earth. Sensory stimuli from our surroundings flood our senses every single day, and we use this information to make meaningful actions and interact with the world. But how does one go, for example, from seeing a traffic light turn green to pushing the gas pedal? The integration of incoming sensory information from the environment with one's prior knowledge and motivations, and the subsequent use of this information to guide actions and decisions is a fundamental task of the nervous system[1–5]. Neural activity, and thus information, flows across the brain from sensory regions to motor regions during movement preparation and execution[6–8], in a time-dependent hierarchical manner through many individual cortical areas[9]. Which cortical areas are directly involved, and the exact mechanisms by which sensory information is converted to motor plans, however, are still being elucidated today.

The Parietal Motor Network

The primary motor cortex (M1) has long been known to be the main movement output center of the brain since its discovery[10] and its subsequent subdivision based on somatotopy[11]. The precise mechanisms by which M1 activity is commands movement are still unknown, and are the subject of much debate in the field today. Activity in M1 during movement has been shown to reflect many different

components of movement such as speed[12], movement direction[13], and control of the muscles involved[14]. However, both M1 and the premotor cortex (M2) have shown preparatory activity in the absence of movement which cancels out population level[15], indicating that both motor regions process information that is distinct from movement output itself, and are more than simply output areas. Rather than encoding a combination of kinematic parameters, recent theories suggest that a more accurate view is that M1 integrates information from multiple different areas to form a motor plan at the population level[16,17]. One of these areas is M2, the secondary motor cortex.

The secondary motor cortex has long been considered essential in goal-directed movement[18], showing strong modulation of activity prior to and during movement[19,20]. In mice, M2 is strongly reciprocally connected to M1 and heavily influences M1 activity[21,22], implying a very strong role of M2 in motor planning. However, the nature of this interaction in mice is still not yet fully described.

The posterior parietal cortex (PPC) is another area that shows very strong modulation to movement preparation and while an animal keeps a decision in mind[23], and projects strongly to M2[23,24]. In monkeys, PPC inactivation has been shown to cause impairment of hand-eye coordination[25,26], specifically impairing reaches but not saccades, implying PPC as being essential in the sensorimotor transformation. Together, these three areas form the parietal motor network, and are the three areas of interest in this study.

How does information flow between these areas during a sensorimotor task in mice? In monkeys, pharmacological inactivation of M1 has been shown to block

electrically-evoked movements from M2 and PPC, while inactivation of PMC blocks electrically-evoked movements from PPC[27], implying an connectivity hierarchy, and information flow, of PPC to M2, and M2 to M1. Taking this into consideration, inactivation of further upstream areas like the PPC may likely have a smaller behavioral effect than inactivation of downstream areas, like M1.

The Mouse as a model

The mouse brain shares many cortical similarities to humans and other higher mammals, but its lissencephalic structure and lower modularity allows for access to large parts of the dorsal cortex for easier study[28,29]. Well-developed transgenic lines established through the use of the Cre-recombinase[30] allow for neuronal cell-type specific targeting of optogenetic tools such as Channelrhodopsin2 to be transduced into all neurons or specific neuronal subtypes, such as parvalbumin(+) GABAergic interneurons[31–34]. This allows for inactivation of cortical with high temporal resolution that cannot be achieved through physical lesion or pharmacological inactivation (e.g. via muscimol[35]), allowing us to assess the requirement of M1, M2, or PPC in a novel sensorimotor discrimination task while avoiding potential cortical compensation[36] from other areas.

Results

A Novel Visual Discrimination Task

We developed a two-alternative forced-choice[1] non-memory visual-discrimination task for head-fixed mice using a joystick setup (Figure 1.1A and S1A). In each trial, mice were cued by drifting gratings either forward or downwards for 2

seconds followed by an auditory go cue, after which they were required to press in the corresponding direction of the stimulus to receive reward. The stimulus was presented throughout the response period of the task, and any movement out of the initial hold zone during the visual cue caused a trial abortion. Mice performed one session per day until. After approximately 2 months of training (Figure S1B), animals performed the task proficiently (Figure 1.1B). Many animals developed a constant target preference within each session, quantified as a preference index (Methods), with no trend towards either target across animals ($p=0.84$, t-test) (Figure 1.1C). Movements made to an animal's preferred target were more stereotyped (Figure 1.1D), in terms of the ballistic angle, the angle at the point of maximum joystick velocity and off target movement, the path distance of movement towards the opposite target (Figure 1.1E).

M1 Inactivation Exacerbates Constant Target Preference

To probe the role of M1 in the task, we transduced Channelrhodopsin-2 (ChR2) into PV neurons in M1 via Cre-dependent AAV (Figure 2A) and trained these mice (N=10) in the task. Once performance reached sufficient levels (Methods), we began inactivation experiments, applying blue light bilaterally to M1 through thinned-skull windows in approximately 15% of trials per session (Methods), alternating control and inactivation sessions daily (Figure 1.2B). Surprisingly, only 2 of 10 animals were significantly immobilized by inactivation compared to control days ($p<0.01$, Wilcoxon one-sided rank sum test), while overall animals were still able to perform the task during inactivation (Figure 1.2B). Hold Performance was not significantly affected across animals (Figure 1.2C). Discrimination performance, however, was significantly decreased in inactivation sessions compared to control

(Figure 1.2D), due to a large increase in the constant target preference (Figure 1.2F). These inactivated movements were in the more stereotyped target direction (Figure 1.2G).

M2 Inactivation Impairs Withholding and Exacerbates Target Preference

Similarly to M1 animals, we transduced ChR2 into PV neurons in the secondary motor cortex (Figure 1.3A) of 8 mice, trained them in the task, and performed inactivation sessions once performance reached sufficient levels. One animal was significantly immobilized by the light, though overall animals were not (Figure 1.3B). Hold performance significantly decreased across animals ($p<0.05$, Wilcoxon one-sided signed rank test) (Figure 1.3B). Discrimination performance was also significantly decreased ($p<0.01$, Wilcoxon one-sided signed rank test), along with the preference index ($p<0.01$, Wilcoxon one-sided signed-rank test).

PPC is not required for the Visually-Guided Task

To probe the role of PPC, we then transduced ChR2 in PPC (Figure 1.4A) in another set of PV-Cre mice (N=5), trained them in the task, and performed inactivation sessions. Animals were not immobilized and were able to perform the task (Figure 1.4B). Unlike M1 inactivation, however, there were no significant changes in hold performance (Figure 1.4C), discrimination performance (Figure 1.4D), or the preference index (Figure 1.4E) within or across animals, indicating PPC was not necessary for task performance.

Discussion

Information travels across the cortex between individual modules in the brain, and is transformed from sensory information to robust, accurate movements. The primary motor cortex (M1), secondary motor cortex (M2), and posterior parietal cortex (PPC), are all areas that have been shown to have strong modulation of activity in preparation of and during movement, making these areas of strong interest. In this study, we used cutting-edge genetic technologies in combination with a complex behavior task to test the requirement of each of these cortical areas via acute optogenetic silencing, allowing for within-session comparison of behaviors.

The Requirement of M1 for Movement

The exact role of M1 in motor skill execution has been a long-standing question in neuroscience. M1 is known to be essential for making skilled movements[16], finding evidence in lesion experiments in which animals lose the ability to grasp[37]. M1 seems to also be required for motor skill learning, as lesioned animals are incapable of learning new contingencies[38,39]. Because motor learning is defined by stereotypy of movement[38,40], M1 must be required for development of stereotypy of movement. In this study, training in our task took several months to reach acceptable discrimination levels, and animals typically developed more stereotyped movements to one of two targets. Optogenetic inactivation of M1 did not cause immobilization across mice, as seen in previous studies[41], but rather caused an increase in the movements an animal made to its more stereotyped target direction. M1 was either not, or less, required for the performance of these movements, complementing the results of physical lesion studies[39]. It could be that lower brain areas in charge of

motor control, such as the striatum or other parts of the basal ganglia, were taking control in commanding this movement in reaction to a loss of feedforward activity from acute M1 inhibition. It is important to note that our task was a lever-press task and does not necessarily require the animal to grip and maneuver the lever with precision. Rather, the animal only needs to impart force in its choice direction, which only requires large muscles to move and not fine motor control. If the task had required grasping[37,41], licking[7,42,43], or other motor skill, we may have seen immobilization instead.

The Role of M2 in Mice

M2 inactivation in our task produced very similar results to M1 inactivation, but also impaired hold performance in our mice. This result has been seen in rat premotor cortex pharmacological inactivation, which led to impairment of temporal preparation and causing premature responding in reaction time tasks[44,45]. In rats[46,47] and monkeys[15,48,49], premotor areas show strong involvement in perceptual decision-making. It is likely that mouse M2 plays a similar but potentially simpler role in coordinating the temporal specificity of movements.

The strong increase in preference index coupled with a decrease in discrimination analogous to M1 inactivation is interesting, in the least. Stereotypy analysis of the movements was not done, but would have revealed whether the effect was due to subcortical control, or due to M2 activity modulating M1 and to encode choice. The latter phenomenon has been seen in monkeys, where premotor cortical activity tuned to a selected target while primary motor cortical activity to the opposite target was surpressed[49]. This would negate the previous claim that mouse M2 is a simple region, rather supporting that a complex combination of choice, temporal

specificity, and potentially other aspects of motor control are processed in mouse M2. This is supported in a rat study, which showed M2 activity integrating reward responses without interfering with motor output[50].

The discrimination deficit caused by inactivation should be interpreted with caution. As with all optogenetic or pharmacological lesion studies, it could be that a decrease in performance is not necessarily caused by the region directly involved being inhibited, but rather by a lack of feedback activity to connected areas[51]. Because mouse M1 and M2 run very close to each other anatomically[52,53], and are very strongly reciprocally connected[21,22], it could be the sudden loss of feedforward activity from M2 to M1 causes a strong perturbation in M1 activity, leading to simultaneous perturbation of both areas. In fact, one of the eight animals used in the inactivation experiments was significantly immobilized in the task, a phenotype atypical to premotor cortical inactivation, supporting the prior claim. However, this could also be a result of an off target inactivation, an unlikely, but possible concern.

PPC is Nonessential for Task Performance

It seems PPC is dispensable during the entire active phase of our task. No behavioral effect from PPC inhibition has been seen previously in whisker-detection task[7], auditory accumulation of evidence[54], or visually-guided maze task [23], but visually-guided tasks with memory period have shown to require PPC via optogenetic[42] and pharmacological[23] inactivation. Considering the primary visual cortex in mice does not directly project to the forelimb area of M1, only indirectly through PPC[55], it seems strange that PPC inactivation does not cause a significant deficit in visual discrimination, though potentially increasing the sample size may make this difference statistically significant, but small nonetheless. With previous

studies showing PPC deficits most often in tasks with memory periods, it would be worthwhile to test if the addition of a memory period changes the requirement of PPC for this task. Because PPC is known to be heavily involved in decision making[1,2], it could be that inhibition of PPC produces a more subtle change undetectable when looking solely at performance. PPC has been also shown to be essential to hand-eye coordination in monkeys[25,56,57], as such it would be interesting to analyze potential deficits of PPC inhibition in a more complex behavior, such as a grasp task[41] with mice.

Although the results of these perturbations offer some insight into the roles of the mouse parietal motor network, there are obvious drawbacks to optogenetic experiments such as this. As noted earlier, an effect caused by the optogenetic inactivation of a cortical area alone does not necessarily causally link that area to the aspect of deficit, as inactivation can cause off-target effects to projection areas[51].

A lack of quantification of the extent of optogenetic inhibition through the thinned-skull preparation adds unwelcome uncertainty to the specificity of inactivation. A recent study using a craniotomy and glass window in VGAT-ChR2 mice found reduction of activity-dependent marker c-Fos only a few hundred microns from the window edge[42], which would likely be lower in our thinned-skull preparation due to the lower transparency of the skull compared to glass, giving us some confidence in the spatial range of inactivation. It is more likely that we were inactivating a smaller range than intended rather than larger.

Overall, these experiments add much to the growing literature elucidating the identities and exact roles of the cortical areas involved in the sensorimotor transformation. Combined with incredible techniques such as *in-vivo* two-photon

calcium imaging[58,59] to image hundreds of neurons in real-time, the secrets of the brain are more available than ever before.

Acknowledgements

Jeff Dahlen created the stereotypy of movement plots in Figure 1.1D and E, Figure 1.2G and corresponding part of results in Chapter 1.

Figure 1.1. Visual Discrimination Task.

(A) Top: task schematic. Middle: task trial structure. Bottom: stimulus-action contingency.

(B) Withholding (0.846, median) and Discrimination (0.812, median) performance across all mice (n=23 animals, across 245 sessions).

(C) Forward preference index for each session across all mice (0.48 ± 0.18, mean ± std.). There is no target preference across animals (t-test for mean different from 0.5, p=0.84).

(D) Top: example traces to both targets from one session. Animals developed more stereotyped movements to one direction. Bottom: ballistic angle and off-target movement schematic.

(E) Top: Difference in ballistic angle variability vs. preference index. Preferred target movements have lower ballistic angle variability (Pearson's correlation = -0.39, $p<0.005$) Bottom: Difference in off-target movement vs. preference index. Preferred target movements have less off-target movement (Pearson's correlation = -0.51, $p<0.001$; colors, animals; points, individual sessions).

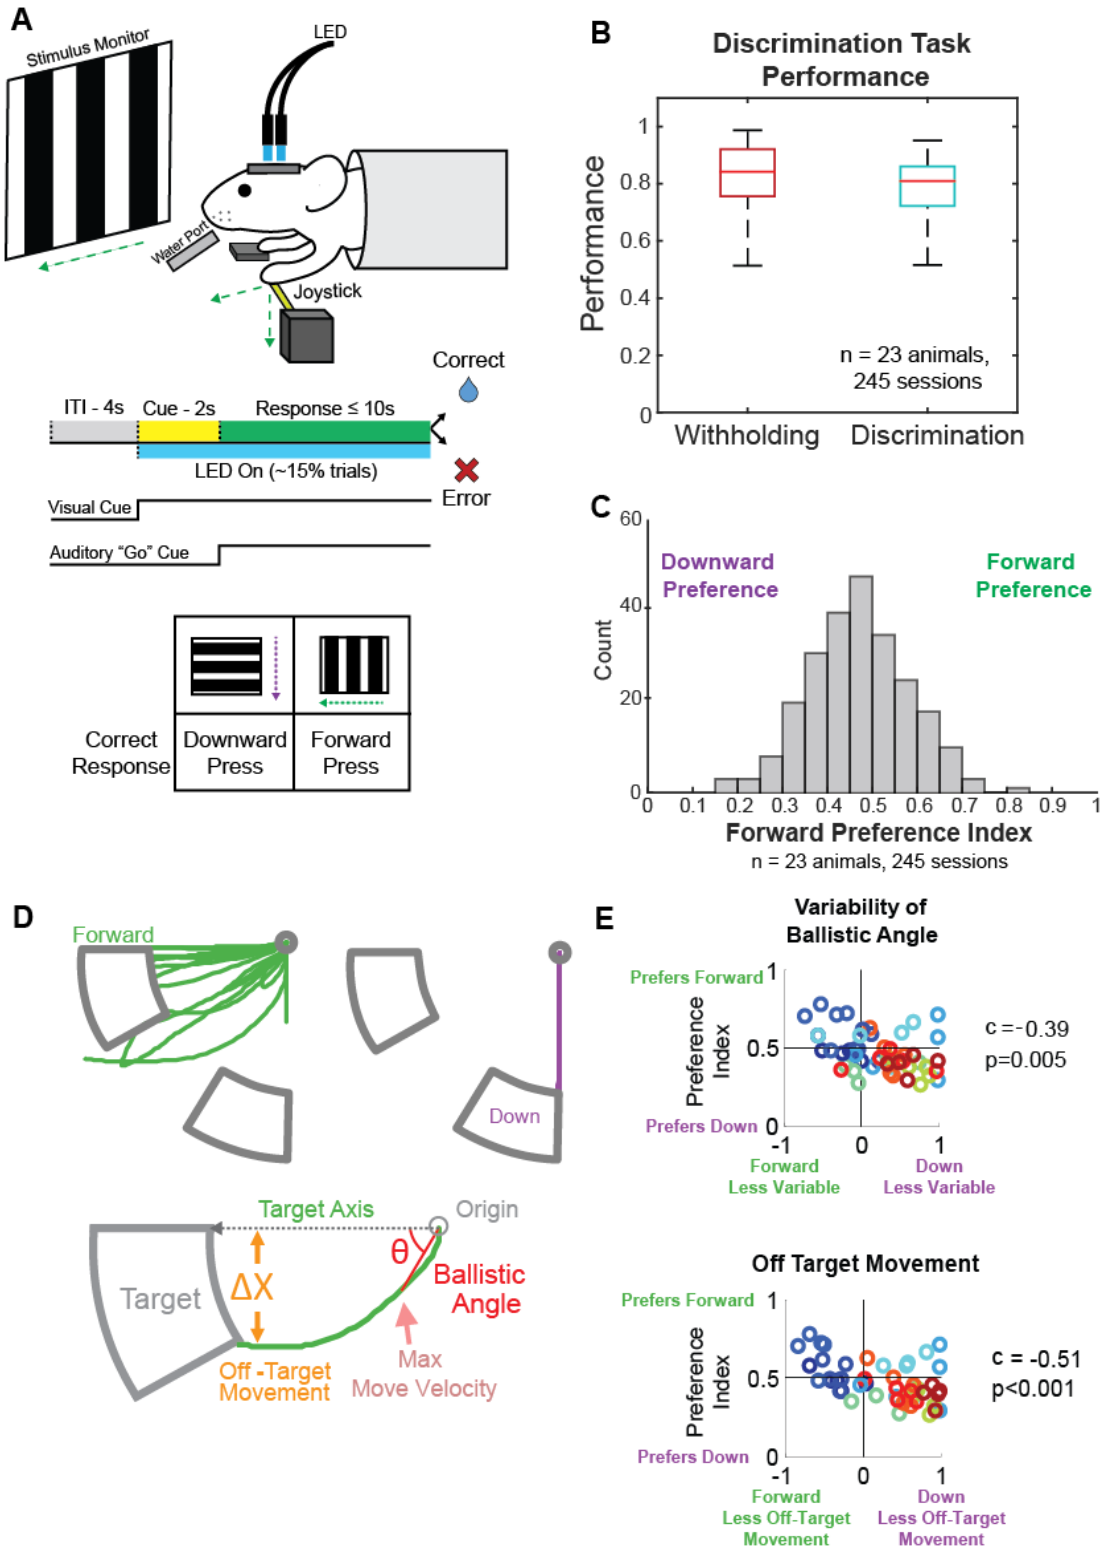
A
Stimulus Monitor
LED
Water Port
Joystick
Correct
ITI - 4s
Cue - 2s
Response ≤ 10s
LED On (~15% trials)
Error
Visual Cue
Auditory "Go" Cue
Correct Response
Downward Press
Forward Press
B
Discrimination Task Performance
Performance
n = 23 animals, 245 sessions
Withholding
Discrimination
C
Downward Preference
Forward Preference
Count
Forward Preference Index
n = 23 animals, 245 sessions
D
Forward
Down
Target Axis
Origin
Target
ΔX
θ
Ballistic Angle
Off -Target Movement
Max Move Velocity
E
Variability of Ballistic Angle
Prefers Forward
Prefers Down
Preference Index
c =-0.39
p=0.005
Forward Less Variable
Down Less Variable
Off Target Movement
c = -0.51
p<0.001
Forward Less Off-Target Movement
Down Less Off-Target Movement

Figure 1.2. Inactivation of the Primary Motor Cortex.

(A) Left: transduction of Channelrhodopsin2-eYFP into the primary motor cortex (Image adapted from mouse brain atlas[53]). Center: ChR2-eYFP expression picture in M1 (scale bar = 500 μm) Right: colocalization of PV and ChR2 (scale bar = 100μm).

(B) Control vs. Inactivation session LED schematic.

(C) Change in the proportion of non-moving trials between light-off and light-on trials between control and inactivation sessions. 2 animals were individually immobilized across sessions ($p<0.01$, Wilcoxon one-sided rank-sum test), but animals were not overall ($p=0.19$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(D) Difference in hold performance between light-off and light-on trials across control and inactivation sessions ($p=0.09$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(E) Difference in discrimination performance between light-off and light-on trials across control and inactivation sessions (-0.01 ± 0.01 vs. -0.21 ± 0.03, mean ± s.e.m.; $p < 0.005$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(F) Change in constant target preference index between light-on and light-off trials across control and inactivation sessions (0.26 ± 0.02 vs. 0.07 ± 0.01, inactivation vs. control, mean ± s.e.m; $p<0.001$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(G) Left: Difference in ballistic angle variability vs. inactivation preference index. Target movements during inactivation have lower ballistic angle variability (Pearson's correlation = -0.45, $p<0.005$) Right: Difference in off-target movement vs. inactivation preference index. Target movements during inactivation have less off-target movement (Pearson's correlation = -0.57, $p<0.001$; colors, animals; points, individual sessions).

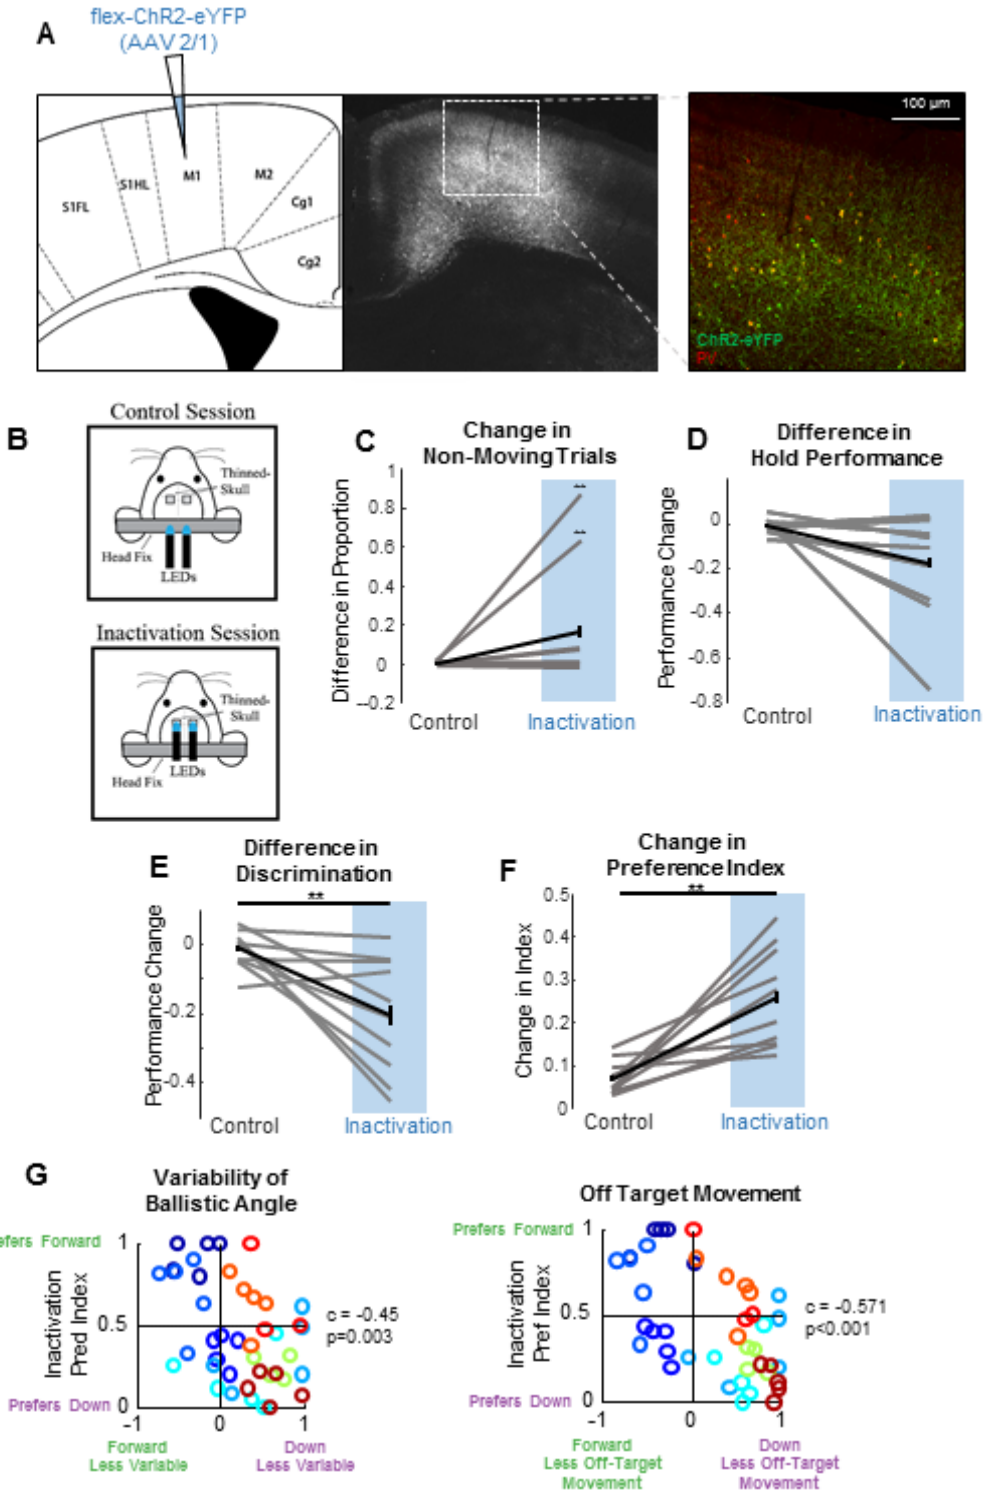
A
flex-ChR2-eYFP
(AAV 2/1)
S1FL
S1HL
M1
M2
Cg1
Cg2
100 µm
ChR2-eYFP
PV
B
Control Session
Thinned-Skull
Head Fix
LEDs
Inactivation Session
Thinned-Skull
Head Fix
LEDs
C
Change in Non-Moving Trials
Difference in Proportion
Control
Inactivation
D
Difference in Hold Performance
Performance Change
Control
Inactivation
E
Difference in Discrimination
Performance Change
Control
Inactivation
F
Change in Preference Index
Change in Index
Control
Inactivation
G
Variability of Ballistic Angle
Prefers Forward
Prefers Down
Inactivation Pred Index
c = -0.45
p=0.003
Forward Less Variable
Down Less Variable
Off Target Movement
Prefers Forward
Prefers Down
Inactivation Pref Index
c = -0.571
p<0.001
Forward Less Off-Target Movement
Down Less Off-Target Movement

Figure 1.3. Inactivation of the Secondary Motor Cortex.

(A) Left: transduction of Channelrhodopsin2-eYFP into the secondary motor cortex. (Image adapted from mouse brain atlas[53]) Right: ChR2-eYFP expression image in M2 (scale bar=500 µm).

(B) Change in the proportion of non-moving trials between light-off and light-on trials between control and inactivation sessions. One animal was significantly immobilized ($p < 0.05$, Wilcoxon one-sided rank-sum test). Animals were not immobilized overall ($p=0.45$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(C) Difference in hold performance between light-off and light-on trials across control and inactivation sessions. (-0.02 ± 0.01 vs. -0.22 ± 0.02; $p<0.05$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(D) Difference in discrimination performance between light-off and light-on trials across control and inactivation sessions (-0.01 ± 0.01 vs. -0.21 ± 0.03, mean ± s.e.m.; $p < 0.005$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(E) Change in constant target preference index between light-on and light-off trials across control and inactivation sessions (0.26 ± 0.02 vs. 0.07 ± 0.01, inactivation vs. control, mean ± s.e.m; $p<0.001$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

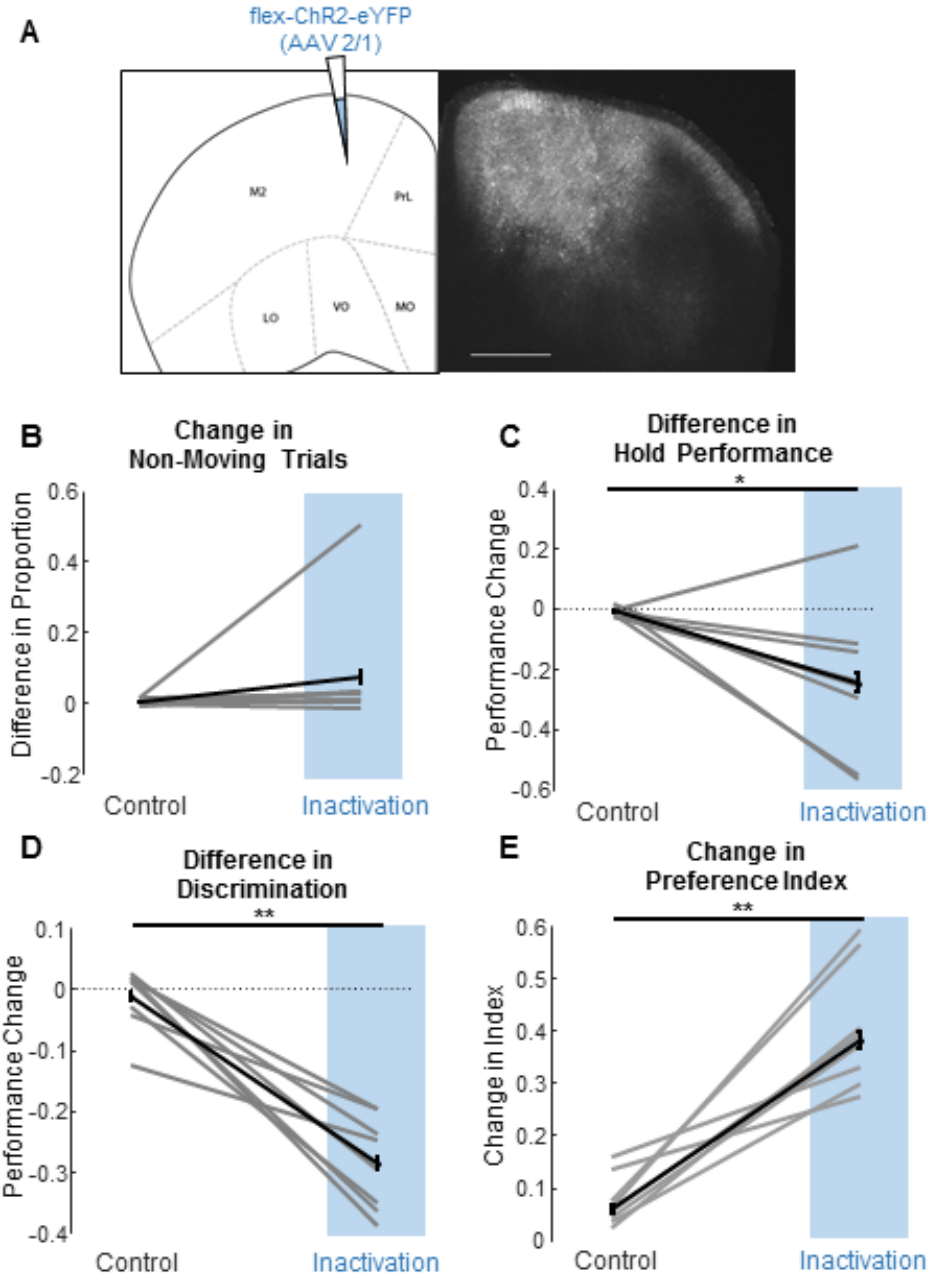
A
flex-ChR2-eYFP
(AAV 2/1)
M2
PrL
LO
VO
MO
B
Change in
Non-Moving Trials
Difference in Proportion
0.6
0.4
0.2
0
-0.2
Control
Inactivation
C
Difference in
Hold Performance
*
Performance Change
0.4
0.2
0
-0.2
-0.4
-0.6
Control
Inactivation
D
Difference in
Discrimination
**
Performance Change
0.1
0
-0.1
-0.2
-0.3
-0.4
Control
Inactivation
E
Change in
Preference Index
**
Change in Index
0.6
0.5
0.4
0.3
0.2
0.1
0
Control
Inactivation

Figure 1.4. Inactivation of the Posterior Parietal Cortex.

(A) Left: transduction of Channelrhodopsin2-eYFP into the posterior parietal cortex. (Image adapted from mouse brain atlas[53]) Right: ChR2-eYFP expression image in PPC (scale bar = 500 μm).

(B) Change in the proportion of non-moving trials between light-off and light-on trials between control and inactivation sessions. Animals were not immobilized overall ($p=0.81$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(C) Difference in hold performance between light-off and light-on trials across control and inactivation sessions. ($p=0.24$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(D) Difference in discrimination performance between light-off and light-on trials across control and inactivation sessions ($p=0.08$, Wilcoxon one-sided signed rank test; individual lines, animals; black, mean ± s.e.m.).

(E) Change in constant target preference index between light-on and light-off trials across control and inactivation sessions ($p=0.12$, Wilcoxon one-sided signed rank test)

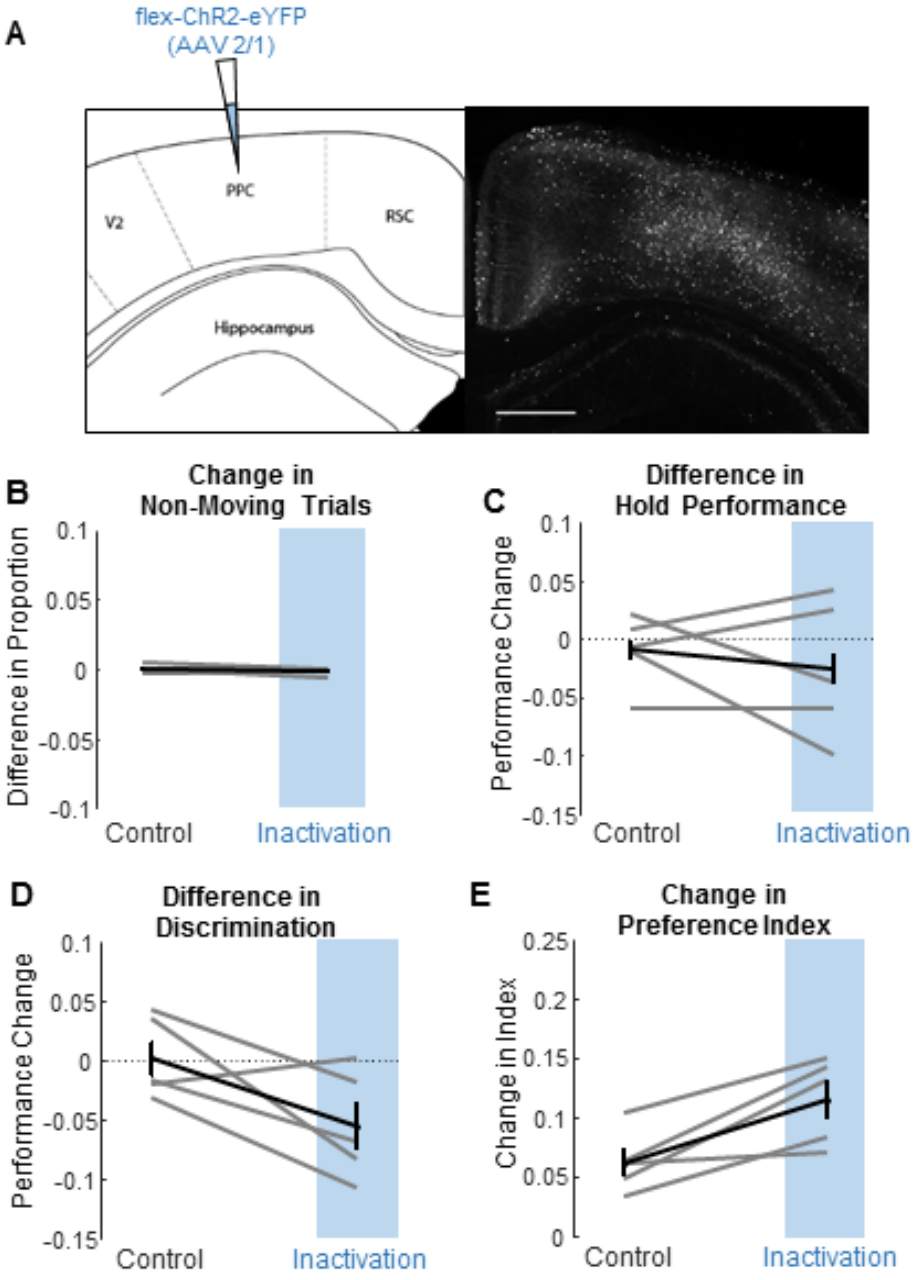
A
flex-ChR2-eYFP
(AAV 2/1)
PPC
V2
RSC
Hippocampus
B
Change in
Non-Moving Trials
Difference in Proportion
0.1
0.05
0
-0.05
-0.1
Control
Inactivation
C
Difference in
Hold Performance
Performance Change
0.1
0.05
0
-0.05
-0.1
-0.15
Control
Inactivation
D
Difference in
Discrimination
Performance Change
0.1
0.05
0
-0.05
-0.1
-0.15
Control
Inactivation
E
Change in
Preference Index
Change in Index
0.25
0.2
0.15
0.1
0.05
0
Control
Inactivation

Chapter 2: Perturbation of PPC alters History-Dependency of Choice

Introduction

Decision Making

As humans, nearly every decision that we make in our adult lives is subject to some sort of bias, that we may or may not even be aware of [60]. These biases are based on the choices that we make, which in turn affect future choices we intend to make[61], and are often subject to change in real-time [46,62,63]. How are these biases created and maintained within the brain? Areas like the prefrontal cortex[48,62,64], premotor cortex[12,24,65], and the parietal cortex[66–69] have shown modulation to choice history and reward history in decision-making tasks. This study focuses on the posterior parietal cortex, and its role in the computation of bias in sensorimotor decision-making.

The Posterior Parietal Cortex

The posterior parietal cortex is known to be heavily involved in decision-making in the sensorimotor pathway[1,2], and activity has been shown to correlate with the subjective desirability of action[70]. Considering neurons even within PPC have very diverse activity profiles that depend heavily on their individual connectivity[23,71], and that PPC activity is known to dynamically change during the decision-making process[72], it is possible that PPC serves as a center for the generation of an internal bias from an analysis of reward and previous choices, and then subsequently communicates this information to its many projections. But when does this happen during the decision-making process? How much sensory information is considered?

While activity in the PPC has been shown to be non-essential in the performance of sensory-guided tasks[7,54], pharmacological inactivation of PPC in mice has been shown to reduce performance of a visually-guided maze task only when a memory period was included, and not when the visual cue was presented throughout the trial[23]. As such, we developed a novel two-alternative forced-choice[1] visually-guided memory discrimination task to study how PPC contributes to the processing of bias in the decision-making process.

In this study, using a combination of mathematical modeling of behavior and well-developed optogenetic technologies[31–34] in mice lines established through the use of the Cre-recombinase[30], we establish a link between pre-stimulus PPC activity and an emergent choice-reward history dependent bias.

Results

Visually-guided Memory Task

We trained mice in a novel sensorimotor memory task (Figure 2.1A and S1A). In this task, mice were cued by drifting gratings either forward or downwards for 1 second, followed by a 2 second memory period during which the mice were required to actively withhold movement. At the end of this period, response was cued via an auditory go cue, after which they were required to press in the corresponding direction of the stimulus to receive reward. Any movement out of the initial hold zone during the visual cue or the memory period caused a trial abortion. Though mice discriminated significantly better than chance, based on an animal's individual constant target preference in an individual session (Figure 2.1B, Methods), performance was significantly worse than in the non-memory discrimination task

(Figure 2.1C), suggesting that memory was the performance-limiting factor rather than visual discriminability.

History-based action selection bias revealed by behavioral modeling

The high choice variability of in the task was not due to a constant target preference, and could have been due to random choice or a much more systematic fluctuation of hidden internal biases, formed from the animal's choice and reward history. To test the latter possibility, we built a logistic regression model of the behavior that predicts the target choice of each mouse for individual trials within a session as a function of the sensory stimulus direction of the current trial, the choice-outcome history from previous trials, and a constant choice preference (Equation 1 and Methods). The portion of the equation excluding the stimulus direction corresponds to the estimate of the internal bias in each trial. Regression was individually performed for each session to determine the weights and time constants of the history terms that best fit the animal's choice. To avoid overfitting, model accuracy was computed via cross-validation in which the model was built using a fraction of the trials in a session ('training set') and evaluated for the accuracy on the remaining trials ('test set').

$$\log \frac{probability\{choice(N)=forward\}}{1-probability\{choice(N)=forward\}} = w_s \cdot stimulus(N) + w_o \cdot \sum_{k=1}^{N-1} outcome(k) \cdot e^{-\frac{N-1-k}{\tau_o}} + w_c \cdot \sum_{k=1}^{N-1} choice(k) \cdot e^{-\frac{N-1-k}{\tau_c}} + w_{oc} \cdot \sum_{k=1}^{N-1} outcome(k) \cdot choice(k) \cdot e^{-\frac{N-1-k}{\tau_{oc}}} + constant \quad (1)$$

Strikingly, this model with choice-outcome history predicted the behavior significantly better than the stimulus alone (Figure 2.2A and 2.2B), indicating that the choice variability in this task was not random. Instead, a significant part of the variability arose from a systematic influence of choice-outcome history that biased the

decision on a trial-by-trial basis. Model accuracy was significantly worse when a model built for one session was applied to another session of the same mouse, and was even lower when a model from one animal was applied to another mouse, demonstrating an idiosyncratic nature of the strategies (Figure 2.2C). Taken together, the imperfect behavioral performance in conjunction with the behavioral model gives us an opportunity to estimate the hidden internal biases underlying decision variability in individual trials which are not directly measurable.

Optogenetic perturbation of PPC ITI activity alters internal bias

To test the involvement of PPC in the determination of this internal bias, we used optogenetics to perturb PPC activity during the task. We transduced ChR2 into PPC neurons in mice and bilaterally perturbed PPC activity via blue light during the inter-trial interval (ITI) of a small subset (~15%) of trials (Figure 2.3A). If PPC is indeed essential for the decision biases based on choice-outcome history, then PPC perturbation should alter the idiosyncratic relationship between history and choice. We tested this idea by building a behavioral model with a subset of the light-off, unperturbed trials as the training set, and testing the accuracy of choice prediction for the remaining unperturbed and perturbed trials respectively. Strikingly, the model built with unperturbed trials was significantly better at predicting the choice of other unperturbed trials than for perturbed trials (Figure 2.3B-D). This result shows that PPC perturbation altered the idiosyncratic relationship between choice-outcome history and the subsequent actions. This effect was not observed in control sessions of the same mice, in which the light was directed at the head bar instead of PPC, indicating that the effect was due to PPC perturbation and not due to non-specific effects of the light (Figure 2.3D).

In a separate set of experiments, we transduced ChR2 in PV-Cre mice to express ChR2 in PV inhibitory neurons in PPC. We performed analogous experiments with these mice in which PPC was inactivated bilaterally in ~15% of trials. Similarly to PPC activation described above, inactivation also altered the history-choice relationship specifically in inactivated trials, further supporting that PPC ITI activity is essential for the history-based action biases (Figure 2.3E). Intriguingly, however, the history-choice relationship was not altered when we inactivated PPC from the beginning of the visual stimulus until the end of the trial (Figure 2.3F and 2.3G), indicating that stimulus, delay and movement responses in PPC are dispensable for the performance of this task.

Discussion

When making decisions, our choices are often strongly affected by personal biases based on the previous decisions we've made [61,63]. In this study, the difficulty mice had in performing the memory task resulted in animals relying on a reward-choice history-based bias that varied session to session to make decisions in the task, causing them to stray from optimal performance. Through robust mathematical modeling, and the use of advanced optogenetic tools in mice, we were able to show that the PPC plays a key role in mediating this bias.

Emergence of History-dependent Bias

Previous studies have shown when sensory evidence is difficult to parse, or in this case to remember, that animals increasingly rely on a reward-history based bias[69,73,74] when making decisions. These biases can be modeled mathematically, showing trial to trial variability in choice may seem random, but in actuality is not[2,74].

Here, our logistic model revealed the complex relationship between individual animals' choice-reward history and its directional choices made in the task. Importantly, bias was also shown to be unique in every session, even within animals, indicating that it was actively created during performance of the task. In this way, the reward-choice bias reflective of the highly variable, situational nature of human decision-making[63,64].

The Function of the Posterior Parietal Cortex

This study found the PPC to non-essential for the sensorimotor transformation occurring during the active portion of the memory task[23,54,75], so bias information encoded in the PPC during the ITI must have been transferred to downstream areas before stimulus onset, and these areas not have required constitutive activity in PPC. Rather than being directly involved in movement planning, it seems that PPC is instead is important for mediating subjective decisions[54,72] in this task. Activity in PPC has previously been shown to dynamically change as an animal is actively involved in decision-making, supporting the notion that this history-based bias is calculated in real-time[72], and supported by our model. Neurons in the PPC are known be widely connected within the brain, with the connectivity of individual neurons greatly affecting their activity[23,55,71]. Therefore, the next step to take should be identification and analysis of the activity of the many areas PPC projects to. One of these areas is the prefrontal cortex, which has shown activity correlations with analysis of value[3], previous choices[49,65], and reward associations[62], and is known to be actively involved in movement itself. It would be interesting to see how information from the PPC is integrated into activity in these regions prior to movement execution.

It is important to note that trial inactivation of PPC could have produced a more subtle change undetectable when looking choice and performance. PPC is known to be essential in hand-eye coordination[25,56,57], so trial inactivation could have caused a subtle effect in fine motor skill that was not detectable in our task structure. To further probe the role of PPC, it would be interesting to determine the potential deficits caused by inactivation during a reach or grasp task rather than the decision-making paradigm used here.

A lack of quantification of the extent of optogenetic perturbation through the thinned-skull preparation unfortunately clouds certainty of spatial specificity. A recent study using a craniotomy and glass window in VGAT-ChR2 mice found reduction of activity-dependent marker c-Fos only a few hundred microns from the window edge[42]. In our study, this spread is more likely smaller due to the use of a thinned-skull preparation, giving us some confidence in the spatial extent of perturbation.

Overall, this study was able to show that PPC activity during the ITI was responsible for the computation of a reward-choice history dependent bias that strongly affected mouse behavior in our task. Slowly, we are starting to reveal the neural bases for individuality and understanding the complex unconscious computations made by the brain during the decision-making process.

Acknowledgements

The results section, as well as Figures 2.1, 2.2, 2.3, and S1 of this chapter are adapted from material that that has been submitted for publication. Hwang, EunJung; Dahlen, Jeff E.; Mukundan, Madan; Komiyama, Takaki. “History-based Action

Selection Bias in Posterior Parietal Cortex". The thesis author was a co-author of this work.

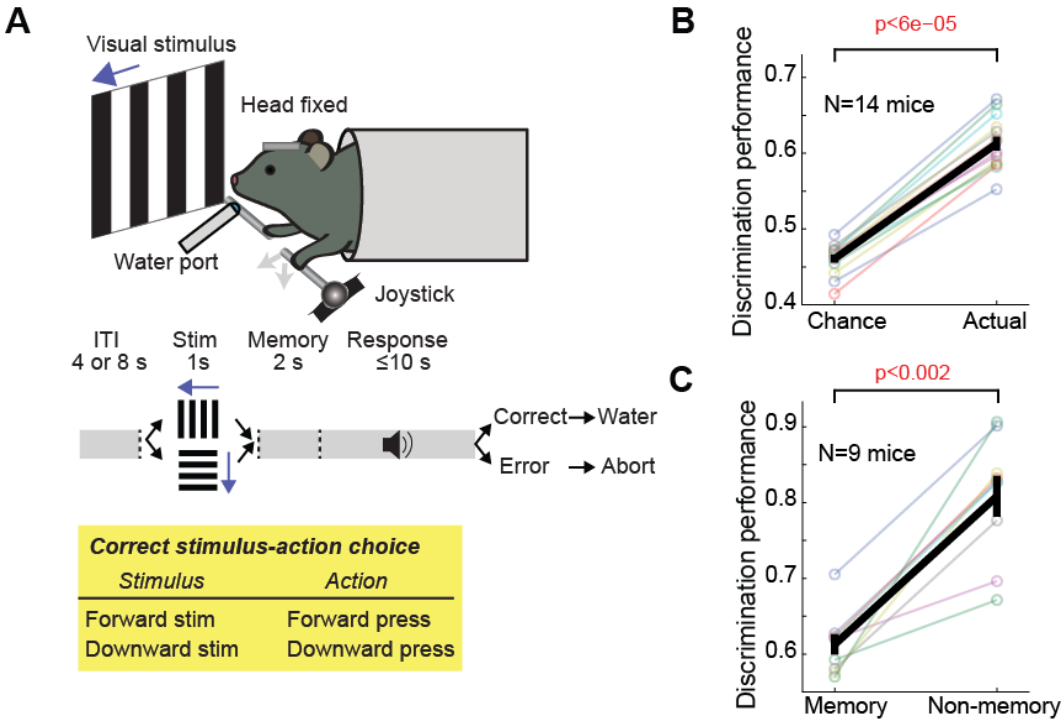

Figure 2.1. Visually-guided memory task.

(A) Top: task schematic. Middle: task trial structure. Bottom: stimulus-action-outcome rule.

(B) Discrimination performance is significantly better than chance. (Black, mean ± s.e.m. across mice; colors, individual mice. Wilcoxon one-sided signed rank test.)

(C) Discrimination performance is significantly better in non-memory trials (visual stimulus stays on throughout the trial) when randomly interleaved with memory trials. (Black, mean ± s.e.m. across mice; colors, individual mice. Wilcoxon one-sided signed rank test.)

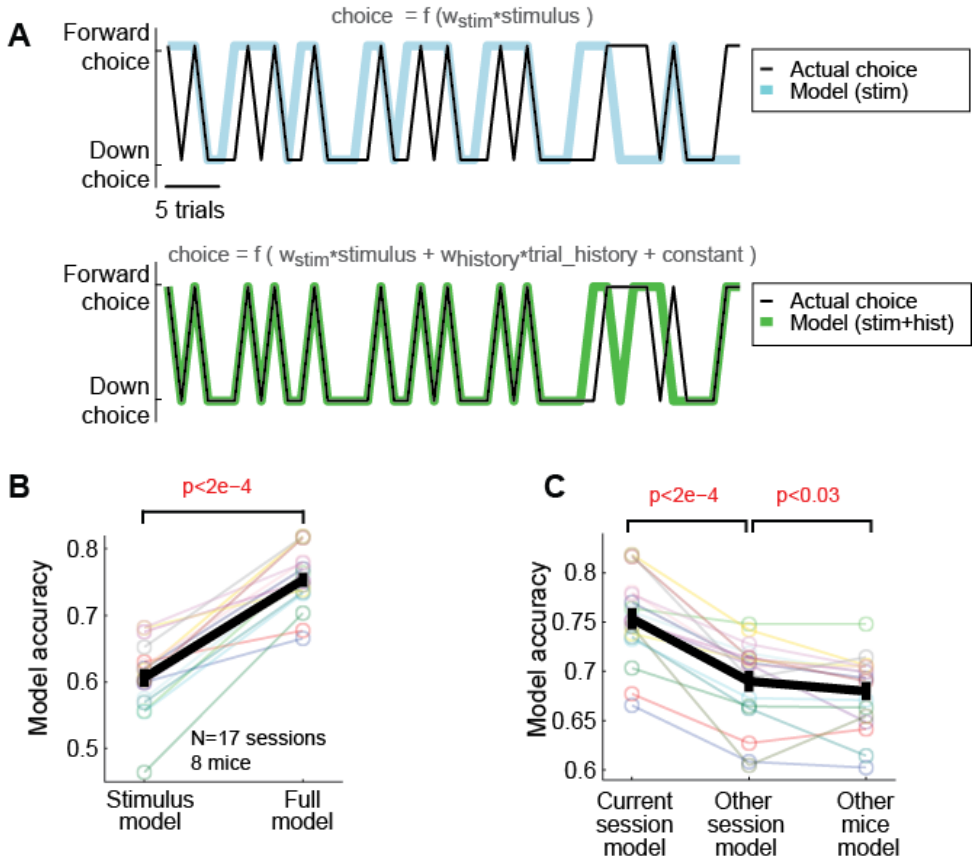

Figure 2.2. Choice-outcome history biases future decision, driving choice variability.

(A) Top: example trial-by-trial sequence of an animal's choice (black) with partial model fit using stimulus information only (cyan). Bottom: choice sequence from top plot (black) with full model fit using stimulus, trial history, and a constant (green).

(B) Full model including both stimulus and history information predicts choice more accurately than partial model including only stimulus. (Black, mean ± s.e.m. across sessions; colors, individual sessions. Wilcoxon one-sided signed rank test.)

(C) Choice prediction accuracy is highest when models are estimated from the current session than different sessions of the same animal, or different animals, indicating that history-dependent strategies vary across sessions and animals. Black, mean ± s.e.m. across sessions; colors, individual sessions. Wilcoxon one-sided signed rank test.

Figure 2.3. Perturbing pre-stimulus activity in PPC alters internal bias and decision performance.

(A) Schematic of perturbation experiment. Control (blue light directed away from PPC) and perturbation (the light directed to PPC) sessions alternated day-to-day (for 14-16 days). Continuous blue light was applied during the ITI in randomly selected trials (15%) in both control and perturbation sessions.

(B) Choice sequence (black) and behavioral model fit (orange) in an example perturbation session. In this example, the mouse tended to alternate choice (i.e., the mouse most heavily weighted the previous choice history) in light-off trials, but this tendency was reduced in light-on trials.

(C) The effect of PPC ITI activation on the model fit in 8 separate activation sessions in a single mouse. (Black, mean ± s.e.m. across sessions; colors, individual sessions. Wilcoxon one-sided signed rank test.)

(D) Average model accuracy in light-off and light-on trials in activation versus control sessions. (Black, mean ± s.e.m. across mice; colors, individual mice. Wilcoxon one-sided signed rank test.)

(E) The same as **D**, but for ITI inactivation.

(F) In post-stimulus-onset inactivation sessions, blue light was applied from stimulus onset to the end of randomly selected trials (15%).

(G) The same as **D**, but for post-stimulus-onset inactivation.

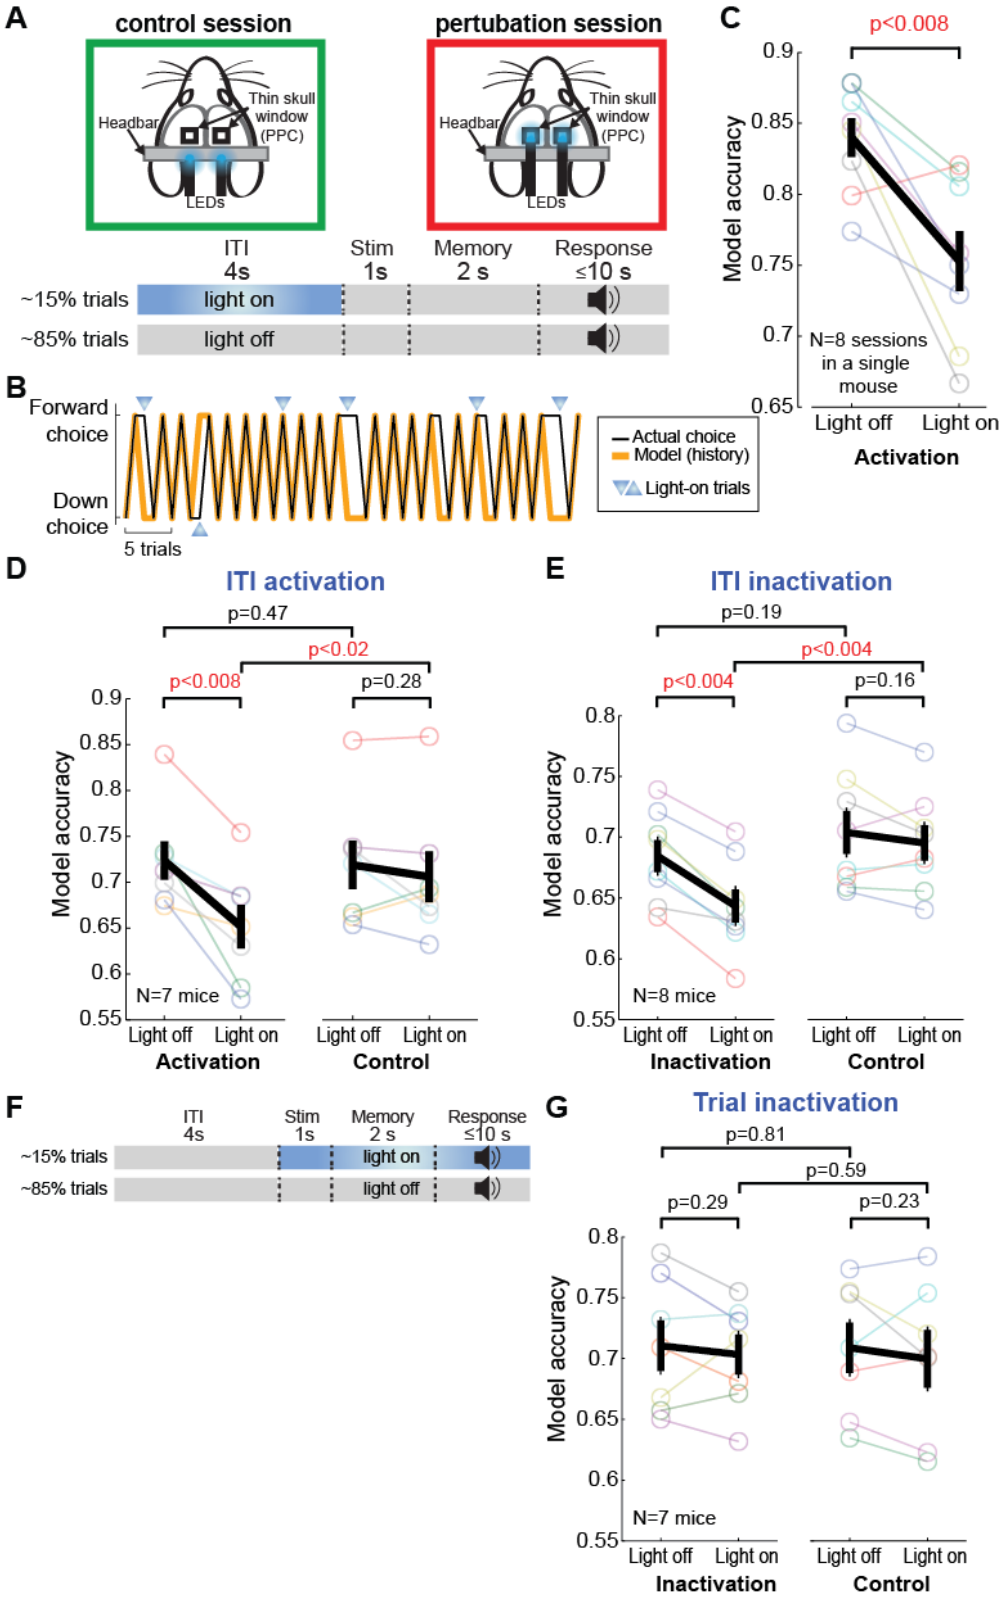
A
control session
pertubation session
Thin skull window (PPC)
Headbar
LEDs
ITI 4s
Stim 1s
Memory 2 s
Response ≤10 s
~15% trials
light on
~85% trials
light off
B
Forward choice
Down choice
5 trials
Actual choice
Model (history)
Light-on trials
C
p<0.008
Model accuracy
N=8 sessions in a single mouse
Light off
Light on
Activation
D
ITI activation
p=0.47
p<0.02
p<0.008
p=0.28
N=7 mice
Activation
Control
E
ITI inactivation
p=0.19
p<0.004
p<0.004
p=0.16
N=8 mice
Inactivation
Control
F
G
Trial inactivation
p=0.81
p=0.59
p=0.29
p=0.23
N=7 mice

Methods

Animals

All experiments were performed in accordance with protocols approved by the UCSD Institutional Animal Care and Use Committee. PV-Cre [JAX 008069][76] mice were group housed in a room with a reversed light cycle (12h – 24h). Experiments were performed during the dark period.

Surgery

Adult mice (6 weeks or older, male and female) were anaesthetized with 2% isoflurane and given dexamethasone (anti-inflammatory, 2 mg/kg), baytril (antibiotic, 10 mg/kg), and buprenorphine (analgesic, 0.1 mg/kg) subcutaneously at the beginning of surgery. The scalp was cleaned with betadine, then removed from between the ears to between the eyes. A head-fixation bar was glued to the skull using cyanoacrylate glue, and cemented in position using dental acrylic. A virus carrying Channelrhodopsin-2 (AAV2-1-CAG-ChR2 for activation, AAV1.EF1A.DIO.hChR2 for inactivation, undiluted; UPenn Vector Core) was injected bilaterally through an ~1 mm by ~1 mm thinned-skull window either in M1 (0.3 mm anterior, 1.5 mm bilateral from bregma)[38,77], M2 (2.5 mm anterior, 0.8 mm bilateral from bregma)[22,78], or PPC (2 mm posterior, 1.7 mm bilateral from bregma)[23,42] either in the center of the window at both 200 µm and 600 µm (~100 µL each) in depth from the dura, or in 2-4 locations (~100 µL each) near the center of the window 250 µm in depth from the bottom of the window. The skull was then sealed with a layer of cyanoacrylate glue.

Perfusion and Histology

Animals were anaesthetized with ketamine and perfused with 4% paraformaldehyde. The brain was removed and placed in 4% PFA for 24 hours, followed by at least 24 hours in 30% sucrose solution. Tissue was sectioned coronally at 30 µm. For all staining, sections were blocked in 3% normal goat serum, 0.3% bovine serum albumin, and 0.2% Triton X-100 in PBS for 1 hour at room temperature. They were then incubated in primary antibody solution for 24 hours at 4ºC (chicken Anti-GFP 1:1000, GFP-1020, Aves Labs, Inc.) for GFP staining. Sections were washed 3 times for 10 minutes each in PBS, then incubated for 2 hours in 0.3% bovine serum albumin, and 0.2% Triton X-100 in PBS secondary antibody solution (AlexaFluor488 goat Anti-chicken 1:1000, Invitrogen). Sections were washed again in PBS, then mounted and cover-slipped. Sections were imaged on either a Zeiss AxioZoom.V16 wide-field microscope or Zeiss Imager M2 with Apotome2 system.

Behavioral Training

Adult mice (6 weeks or older, male and female) were allowed to recover for at least 3 days after surgery, after which their daily water intake was restricted to approximately 1 mL/day. Behavioral training started after at least 3 days of water restriction.

A custom behavioral apparatus in a Skinner box (40cm x 40cm x 40 cm) included a joystick (M11L061P; CHProducts), a visual stimulus presentation monitor (17”, placed ~15 cm from the right eye of the mouse), and a water port (Figure 1A). The joystick, with a dynamic range of 56˚ in each angular direction on a spherical surface, was positioned such that mice used their left paw to manipulate it (Fig. S1a).

The 2D angular position of the joystick was continuously monitored and recorded at 1 kHz using a data acquisition card (USB6008; National Instruments) and a custom Matlab program. The task-sequence execution, stimulus selection, auditory cue presentation, reward dispensation, and task time recording were handled by an open source realtime Linux/Matlab software package BControl (http://brodywiki.princeton.edu/bcontrol/). The presentation of visual stimuli (100% contrast, full-field, square wave drifting gratings 0.04 cycles/degree and 3 cycles/sec) was implemented using an open source Matlab toolbox Psychtoolbox (http://psychtoolbox.org/).

Head-fixed mice[29] were trained in the behavioral apparatus, approximately 1 hour per day over a period of 2-4 months. The task gradually evolved to become more complex through multiple steps (Fig. S1b). In the first step, the mice received a water reward if they moved the joystick to the correct target within a 30-second response period even if they hit the incorrect target first. As they learned reliably move the joystick in both directions for reward, we increased the target distance up to 11.1˚ (~10 mm) and decreased the response time to 10 seconds until they could reach distant targets within 10 seconds in more than 80% of trials (step 2). In steps 1 and 2, the joystick was mechanically immovable until the auditory go cue. Starting in step 3, to prevent mice from pushing or leaning onto the joystick before the go cue, we enabled the joystick from stimulus onset and rewarded mice only if they moved the joystick after the go cue and reached the correct target. Error trials in which mice responded before the go cue were aborted with a white noise sound. Step 3 of training continued until achieving withholding performance above 80%.

$$Witholding\ performance = \frac{\#\ of\ responding\ trials\ after\ go\ cue}{\#\ of\ all\ responding\ trials}$$

In step 4, mice were trained to directly reach for the correct target after the go cue by aborting trials as soon as they hit the incorrect target or moved before the go cue. Step 4 continued until they achieved both discrimination and withholding performance above 80%.

$$Discrimination\ performance = \frac{\#\ of\ trials\ hitting\ the\ correct\ target}{\#\ of\ trials\ hitting\ any\ target}$$

Discrimination performance was computed for all trials that reached a target regardless of whether or not the trials were successfully withheld. Once performance reached these levels, the ITI length gradually increased up to 4 seconds (step 5). The task at this point is the non-memory visual discrimination task presented in Figure 1 and used in the inactivation experiments described in Chapter 1 and Figures 1-4.

In step 6, we turned off the visual stimulus simultaneously with the auditory go cue. In step 7, the stimulus period was shortened to 1.8 sec and a 0.2 sec memory period was introduced. In the final step, the stimulus period was gradually decreased down to 1 s and the memory period was gradually increased up to 2 s. As the memory period became longer, the discrimination performance deteriorated and rarely improved above 60% even after a long period of training, so we trained each mouse until their discrimination performance in the 2-sec memory task reached 60% on average.

In discrimination tasks (from step 4 onward), the visual stimulus was randomly selected between forward or downward drifting gratings with the following exceptions: 1) after three consecutive successful trials in one direction, the stimulus always

switched to the other, and 2) after error trials, the same stimulus was repeated. These exceptions were implemented to discourage the mice from developing a strong bias of choosing only one direction and settling at 50% discrimination accuracy. The visual stimulus was randomly selected between forward or downward drifting gratings with the following constraints: 1) after three consecutive rewarded trials in one direction, the stimulus always switched to the other direction, and 2) after error trials, the same stimulus was repeated. These constraints were implemented to discourage the mice from choosing only one direction and settling at 50% discrimination accuracy. Despite the deterministic stimulus after an error or a third consecutive reward in one direction, the mice performed only slightly better in those trials than random trials (Wilcoxon one-side signed rank, $p<0.04$; Figure S1C), indicating that mice did not fully utilize these hidden stimulus presentation rules to their advantage.

Because of this pseudo-random rule in stimulus presentation, discrimination performance achieved by random choice would not be 50% if there was a constant target preference, so we estimated the constant target preference within a session and converted it to a probability to choose each choice using the following formulae:

$$\text{Probability of choice 1} = \frac{1}{2} \times \left(\frac{\#\ of\ trials_{choice\ 1\ |\ stimulus\ 1}}{\#\ of\ trials_{stimulus\ 1}} + \frac{\#\ of\ trials_{choice\ 1\ |\ stimulus\ 2}}{\#\ of\ trials_{stimulus\ 2}} \right)$$

$$\text{Probability of choice 2} = 1 - \text{Probability of choice 1}$$

Then, the chance level performance for the given session was computed by simulating random binary choice with the estimated probabilities under the same pseudo-random rules 1000 times (Figure 2.1B). The constant target preference index metric used in Chapter 1 is equivalent to the probability of choice 1 value.

A subset of mice performed randomly interleaved non-memory and memory trials (both with a 3-sec pre-movement period between the stimulus onset and the go cue; in non-memory trials, the stimulus stayed on until a target was reached) in some sessions during their training period. In those sessions, the discrimination accuracy was consistently lower in memory than non-memory trials, indicating that the memory load, rather than the sensory discriminability, impaired performance in the memory task (Fig. 2.1C).

Perturbation Experiments

Once a mouse reached the 80% discrimination criterion for the non-memory visual discrimination task, or 60% discrimination for the memory task, we conducted 1-7 light acclimation sessions to minimize non-specific light effects on behaviors. In the acclimation sessions, bifurcated blue LED fibers (wave length: 470 nm, ~4.5mW, Doric) were placed ~1 mm above the head-fixation bar, away from the cortical region expressing Channelrhodopsin2, and lights were turned on in randomly selected 15% of trials. Most mice recovered their usual task performance within a day or two.

Each perturbation experiment was performed across 10 or 16 daily sessions. Control and perturbation sessions alternated day-by-day for all but 5 mice doing memory task inactivation. These 5 mice (3 trial and 2 ITI inactivation) performed control and perturbation sessions sequentially in 7-day blocks. In control sessions, the LED lights were directed above the head-fixation bar, whereas they were placed directly above M1, M2, or PPC on both hemispheres in perturbation sessions. Except for this difference, all procedures were identical.

In both control and perturbation sessions, light trials were pseudo-randomly selected with a restriction that there should be at least 5 no-light trials between any two adjacent light trials, to avoid potential behavioral adaptation to cortical perturbation due to consecutive and/or frequent exposures. Under this restriction, light stimulation was applied to less than 15% of trials.

Trial selection for Behavioral Model

In behavioral model analyses, the response choice was predicted only for trials in which mice reached any of the two targets after the stimulus onset, ~248 trials (range: 137-344; 91%) per session.

Behavioral model

In our behavioral model, the choice on a given trial is predicted by a weighted sum of the current stimulus, the history of past trial outcome, choice, and their interaction, and a constant (Equation 1 and 2). Past trials were temporally discounted in an exponentially decaying manner (i.e. stronger effect from more recent trials) with time constants fit independently for each history variable. Stimulus, outcome, and choice were all binary variables with the value of 1 or -1. However, in trials in which mice did not reach a target, choice was zero and outcome was -1 (error).

We repeated the following procedure for a fixed set of time constants (varying from 0.01 to 100 for each history variable), and selected the time constants and weights that produced the highest model accuracy as the best-fit regression parameters. For given time constants, we first found best weights using logistic regression on a training set (Equation 1), and then estimated the choice sequence in a designated test set using the best weights (Equation 2). The two-step process was

10-fold cross validated to avoid overfitting. That is, trials within a session were divided in 10 non-overlapping parts, and each part served as a test set once, while the other nine parts as a training set. The fit of the model (or simply, model accuracy) was measured as the fraction of test trials in which the estimated choice matched the actual choice.

$$\log \frac{probability\{choice(n)=forward\}}{1-probability\{choice(n)=forward\}} = w_s \cdot stimulus(n) + w_r \cdot \sum_{k=1}^{n-1} reward(k) \cdot e^{-\frac{n-1-k}{\tau_r}} + w_c \cdot \sum_{k=1}^{n-1} choice(k) \cdot e^{-\frac{n-1-k}{\tau_c}} + w_{rc} \cdot \sum_{k=1}^{n-1} reward(k) \cdot choice(k) e^{-\frac{n-1-k}{\tau_{rc}}} + constant \quad (1)$$

$$\widehat{choice}(n) = \begin{matrix} 1, & if\ p > 0.5 \\ -1, & otherwise \end{matrix} \} \quad (2)$$

In partial models, we used a subset of variables and performed the same regression procedure. For instance, when estimating the effect of perturbation on trial-history dependency of choice, we used a partial model without the stimulus term and compared the partial model accuracy between light-on and light-off trials.

To assess the statistical significance of history information in predicting future choices, we applied a likelihood ratio test between the full model and a partial model that contains only stimulus and constant terms. We used $p<0.05$ as a significance threshold.

Movement kinematics analysis

Movement onset was defined as the first time at which the joystick velocity exceeded 22.2˚/s (~20 mm/s) continuously for 20 ms and the joystick moved at least 1.3˚ (~1.1 mm) from the origin. The reaction time was measured as the time between the go cue and movement onset, and the movement time as the time between movement onset and the time when the joystick entered any target region.

Statistical Analysis

When testing the hypothesis that the median of one set of samples is less than the median of the other set acquired from the same animals, we used the non-parametric Wilcoxon one-sided signed-rank test. For two sets of samples acquired from different animals, we used the non-parametric Wilcoxon one-sided rank-sum test.

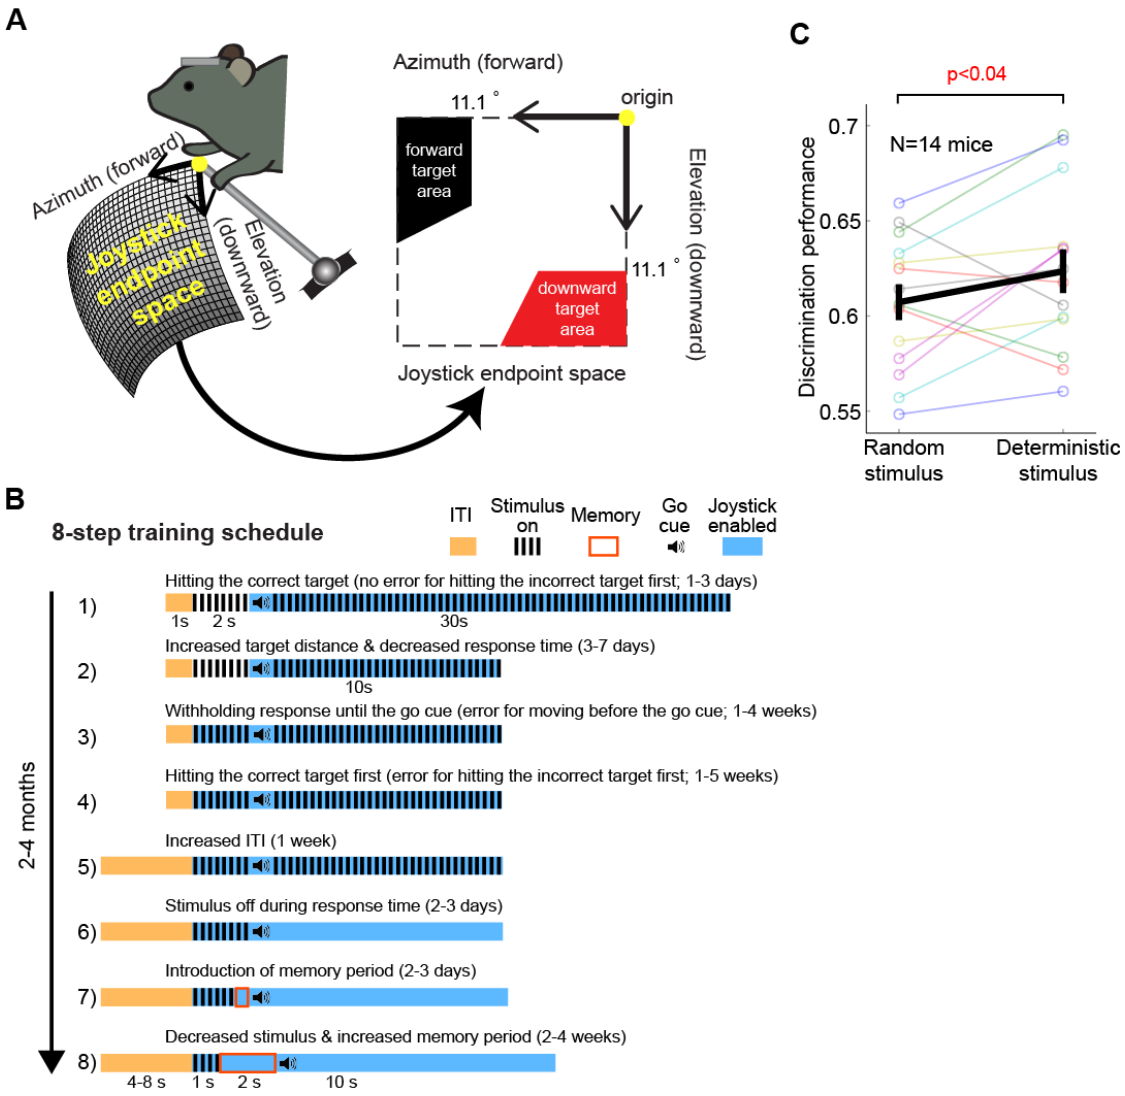

Figure S1. Two-alternative forced-choice tasks using the joystick apparatus.

(A) The two target areas in the joystick endpoint space.

(B) 8-Step behavioral training procedure.

(C) Discrimination performance in random stimulus versus deterministic stimulus trials. Each line represents the average across 20 sessions, colors are individual mice, and the black line is the mean ± s.e.m. across 14 mice. Wilcoxon one-sided signed rank test.

References

1. Carandini M, Churchland AK. Probing perceptual decisions in rodents. *Nat Neurosci*. 2013;16(7):824-831. doi:10.1038/nn.3410.

2. Gold JI, Shadlen MN. The neural basis of decision making. *Annu Rev Neurosci*. 2007;30:535-574. doi:10.1146/annurev.neuro.29.051605.113038.

3. Lee D, Seo H, Jung MW. Neural Basis of Reinforcement Learning and Decision Making. *Annu Rev Neurosci*. 2012;35(1):287-308. doi:10.1146/annurev-neuro-062111-150512.

4. Wang XJ. Decision Making in Recurrent Neuronal Circuits. *Neuron*. 2008;60(2):215-234. doi:10.1016/j.neuron.2008.09.034.

5. Bernacchia A, Seo H, Lee D, Wang X-J. A reservoir of time constants for memory traces in cortical neurons. *Nat Neurosci*. 2011;14(3):366-372. doi:10.1038/nn.2752.

6. Hernández A, Nácher V, Luna R, Zainos A, Lemus L, Alvarez M, Vázquez Y, Camarillo L, Romo R. Decoding a perceptual decision process across cortex. *Neuron*. 2010;66(2):300-314. doi:10.1016/j.neuron.2010.03.031.

7. Guo Z V, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K. Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron*. 2014;81(1):179-194. doi:10.1016/j.neuron.2013.10.020.

8. Romo R, Lemus L, de Lafuente V. Sense, memory, and decision-making in the somatosensory cortical network. *Curr Opin Neurobiol*. 2012;22(6):914-919. doi:10.1016/j.conb.2012.08.002.

9. Murray JD, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Padoa-Schioppa C, Pasternak T, Seo H, Lee D, Wang X-J. A hierarchy of intrinsic timescales across primate cortex. *Nat Neurosci*. 2014;17(12):1661-1663. doi:10.1038/nn.3862.

10. Fritsch G, Hitzig E. Electric excitability of the cerebrum (Uber die elektrische Erregbarkeit des Grosshirns). *Epilepsy Behav*. 2009;15(2):123-130. doi:10.1016/j.yebeh.2009.03.001.

11. Taylor CSR, Gross CG. Twitches Versus Movements: A Story of Motor Cortex. *Neurosci*. 2003;9(5):332-342. doi:10.1177/1073858403257037.

12. Churchland MM, Santhanam G, Shenoy K V. Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *J Neurophysiol*. 2006;96(6):3130-3146. doi:10.1152/jn.00307.2006.

13. Georgopoulos A, Schwartz A, Kettner R. Neuronal population coding of movement direction. *Science (80-)*. 1986;233(4771):1416-1419. doi:10.1126/science.3749885.

14. Townsend BR, Paninski L, Lemon RN, Benjamin R. Linear Encoding of Muscle Activity in Primary Motor Cortex and Cerebellum. *J Neurophysiol*. 2006:2578-

2592. doi:10.1152/jn.01086.2005.

15. Kaufman MT, Churchland MM, Ryu SI, Shenoy K V. Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci*. 2014;17(3):440-448. doi:10.1038/nn.3643.

16. Shenoy K V, Sahani M, Churchland MM. Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci*. 2013;36:337-359. doi:10.1146/annurev-neuro-062111-150509.

17. Michaels JA, Dann B, Scherberger H. Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLOS Comput Biol*. 2016;12(11):e1005175. doi:10.1371/journal.pcbi.1005175.

18. Gremel CM, Costa RM. Premotor cortex is critical for goal-directed actions. *Front Comput Neurosci*. 2013;7(August):110. doi:10.3389/fncom.2013.00110.

19. Li N, Daie K, Svoboda K, Druckmann S, Svoboda K. Robust neuronal dynamics in premotor cortex during motor planning. *Nature*. 2015;532(7600):459-464. doi:10.1038/nature17643.

20. Murakami M, Vicente MI, Costa GM, Mainen ZF. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat Neurosci*. 2014;17(11):1574-1582. doi:10.1038/nn.3826.

21. Hooks BM, Mao T, Gutnisky DA, Yamawaki N, Svoboda K, Shepherd GMG. Organization of Cortical and Thalamic Input to Pyramidal Neurons in Mouse Motor Cortex. *J Neurosci*. 2013;33(2):748-760. doi:10.1523/JNEUROSCI.4338-12.2013.

22. Hira R, Ohkubo F, Tanaka YR, Masamizu Y, Augustine GJ, Kasai H, Matsuzaki M. In vivo optogenetic tracing of functional corticocortical connections between motor forelimb areas. *Front Neural Circuits*. 2013;7(April):55. doi:10.3389/fncir.2013.00055.

23. Harvey CD, Coen P, Tank DW. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*. 2012;484(7392):62-68. doi:10.1038/nature10918.

24. Desmurget M, Sirigu A. A parietal-premotor network for movement intention and motor awareness. *Trends Cogn Sci*. 2009;13(10):411-419. doi:10.1016/j.tics.2009.08.001.

25. Hwang EJ, Hauschild M, Wilke M, Andersen R a. Spatial and Temporal Eye-Hand Coordination Relies on the Parietal Reach Region. *J Neurosci*. 2014;34(38):12884-12892. doi:10.1523/JNEUROSCI.3719-13.2014.

26. Christopoulos VN, Bonaiuto J, Kagan I, Andersen R a. Inactivation of Parietal Reach Region Affects Reaching But Not Saccade Choices in Internally Guided Decisions. *J Neurosci*. 2015;35(33):11719-11728. doi:10.1523/JNEUROSCI.1068-15.2015.

27. Stepniewska I, Gharbawie O a., Burish MJ, Kaas JH. Effects of muscimol inactivations of functional domains in motor, premotor, and posterior parietal cortex on complex movements evoked by electrical stimulation. *J Neurophysiol*. 2014;111(5):1100-1119. doi:10.1152/jn.00491.2013.

28. Krubitzer L. The magnificent compromise: cortical field evolution in mammals. *Neuron*. 2007;56(2):201-208. doi:10.1016/j.neuron.2007.10.002.

29. Guo Z V., Hires SA, Li N, O'Connor DH, Komiyama T, Ophir E, Huber D, Bonardi C, Morandell K, Gutnisky D, Peron S, Xu N, Cox J, Svoboda K. Procedures for Behavioral Experiments in Head-Fixed Mice. *PLoS One*. 2014;9(2):e88678. doi:10.1371/journal.pone.0088678.

30. Van Duyne GD. Cre Recombinase. *Microbiol Spectr*. 2015;3(1):1-19. doi:10.1128/microbiolspec.

31. Madisen L, Mao T, Koch H, Zhuo J, Berenyi A, Hsu YA, Iii AJG, Gu X, Zanella S, Gu H, Mao Y, Hooks BM, Boyden ES, Buzsáki G, Ramirez JM, Jones AR, Svoboda K, Han X, Turner EE. A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat Neurosci*. 2012;15(5):793-802. doi:10.1038/nn.3078.A.

32. Zeng H, Madisen L. *Mouse Transgenic Approaches in Optogenetics*. Vol 196. 1st ed. Elsevier B.V.; 2012. doi:10.1016/B978-0-444-59426-6.00010-0.

33. Suter B a, Yamawaki N, Borges K, Li X, Kiritani T, Hooks BM, Shepherd GMG. Neurophotonics applications to motor cortex research. *Neurophotonics*. 2014;1(1). doi:10.1117/1.NPh.1.1.011008.

34. Taniguchi H, He M, Wu P, Kim S, Paik R, Sugino K, Kvitsani D, Fu Y, Lu J, Lin Y, Miyoshi G, Shima Y, Fishell G, Nelson SB, Huang ZJ. A Resource of Cre Driver Lines for Genetic Targeting of GABAergic Neurons in Cerebral Cortex. *Neuron*. 2011;71(6):995-1013. doi:10.1016/j.neuron.2011.07.026.

35. Hikosaka O, Wurtz RH. Modification of saccadic eye movements by GABA-related substances. I. Effect of muscimol and bicuculline in monkey superior colliculus. *J Neurophysiol*. 1985;53(1):266-291.

36. Kolb B, Brown R, Witt-Lajeunesse A, Gibb R. Neural compensations after lesion of the cerebral cortex. *Neural Plast*. 2001;8(1):1-16. doi:10.1155/NP.2001.1.

37. Whishaw IQ. Loss of the innate cortical engram for action patterns used in skilled reaching and the development of behavioral compensation following motor cortex lesions in the rat. *Neuropharmacology*. 2000;39(5):788-805. doi:10.1016/S0028-3908(99)00259-2.

38. Peters AJ, Chen SX, Komiyama T. Emergence of reproducible spatiotemporal activity during motor learning. *Nature*. 2014;510(7504):263-267. doi:10.1038/nature13235.

39. Kawai R, Markman T, Poddar R, Ko R, Fantana AL, Dhawale AK, Kampff AR, Ölveczky BP. Motor Cortex Is Required for Learning but Not for Executing a

Motor Skill. *Neuron*. 2015;86(3):800-812. doi:10.1016/j.neuron.2015.03.024.

40. Shmuelof L, Krakauer JW, Mazzoni P. How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *J Neurophysiol*. 2012;108(2):578-594. doi:10.1152/jn.00856.2011.

41. Guo J-Z, Graves AR, Guo WW, Zheng J, Lee A, Rodríguez-González J, Li N, Macklin JJ, Phillips JW, Mensh BD, Branson K, Hantman AW. Cortex commands the performance of skilled movement. *Elife*. 2015;4(DECEMBER2015):1-18. doi:10.7554/eLife.10774.

42. Goard MJ, Pho GN, Woodson J, Sur M. Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *Elife*. 2016;5:1-30. doi:10.7554/eLife.13764.

43. Li N, Chen T-W, Guo Z V, Gerfen CR, Svoboda K. A motor cortex circuit for motor planning and movement. *Nature*. 2015;519(7541):51-56. doi:10.1038/nature14178.

44. Narayanan NS, Horst NK, Laubach M. Reversible inactivations of rat medial prefrontal cortex impair the ability to wait for a stimulus. *Neuroscience*. 2006;139(3):865-876. doi:10.1016/j.neuroscience.2005.11.072.

45. Narayanan N, Laubach M. Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. *Neuron*. 2006;52(5):921-931. doi:10.1016/j.neuron.2006.10.021.TOP-DOWN.

46. Sul JH, Jo S, Lee D, Jung MW, Jung Hoon Sul1, Suhyun Jo1, 2, Daeyeol Lee3, and Min Whan Jung1 2. Role of rodent secondary motor cortex in value-based action selection. *Nat Neurosci*. 2011;14(9):1202-1208. doi:10.1038/nn.2881.

47. Sul JH, Kim H, Huh N, Lee D, Jung MW. Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron*. 2010;66(3):449-460. doi:10.1016/j.neuron.2010.03.033.

48. Donahue CH, Lee D. Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex. *Nat Neurosci*. 2015;18(2):1-9. doi:10.1038/nn.3918.

49. Thura D, Cisek P. Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron*. 2014;81(6):1401-1416. doi:10.1016/j.neuron.2014.01.031.

50. Saiki A, Kimura R, Samura T, Fujiwara-Tsukamoto Y, Sakai Y, Isomura Y. Different modulation of common motor information in rat primary and secondary motor cortices. *PLoS One*. 2014;9(6). doi:10.1371/journal.pone.0098662.

51. Otchy TM, Wolff SBE, Rhee JY, Pehlevan C, Kawai R, Kempf A, Gobes SMH, Ölveczky BP. Acute off-target effects of neural circuit manipulations. *Nature*. doi:10.1038/nature16442.

52. Lein ES, Hawrylycz MJ, Ao N, Ayres M, Bensinger A, Bernard A, Boe AF, Boguski MS, Brockway KS, Byrnes EJ, Chen L, Chen L, Chen T-M, Chi Chin M, Chong J, Crook BE, Czaplinska A, Dang CN, Datta S. Genome-wide atlas of gene expression in the adult mouse brain. *Nature*. 2007;445(7124):168-176. doi:10.1038/nature05453.

53. Paxinos G, Franklin KBJ. *Paxinos and Franklin's the Mouse Brain in Stereotaxic Coordinates*. Boston: Elsevier/Academic Press, Amsterdam; 2013.

54. Erlich JC, Brunton BW, Duan C a, Hanks TD, Brody CD. Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife*. 2015;4:1-28. doi:10.7554/eLife.05457.

55. Wang Q, Sporns O, Burkhalter A. Network Analysis of Corticocortical Connections Reveals Ventral and Dorsal Processing Streams in Mouse Visual Cortex. *J Neurosci*. 2012;32(13):4386-4399. doi:10.1523/JNEUROSCI.6063-11.2012.

56. Hwang EJ, Hauschild M, Wilke M, Andersen R a. Inactivation of the parietal reach region causes optic ataxia, impairing reaches but not saccades. *Neuron*. 2012;76(5):1021-1029. doi:10.1016/j.neuron.2012.10.030.

57. Andersen R a, Andersen K, Hwang E, Hauschild M. Optic ataxia: from Balint's syndrome to the parietal reach region. 2012;100(2):130-134. doi:10.1016/j.pestbp.2011.02.012.Investigations.

58. Stosiek C, Garaschuk O, Holthoff K, Konnerth A. In vivo two-photon calcium imaging of neuronal networks. *Proc Natl Acad Sci U S A*. 2003;100(12):7319-7324. doi:10.1073/pnas.1232232100.

59. Akerboom J, Chen T-W, Wardill TJ, Tian L, Marvin JS, Mutlu S, Carreras Caldéron N, Esposti F, Borghuis BG, Sun XR, Gordus A, Orger MB, Portugues R, Engert F, Macklin JJ, Filosa A, Aggarwal A, Kerr RA, Takagi R. Optimization of a GCaMP Calcium Indicator for Neural Activity Imaging. *J Neurosci*. 2012;32(40):13819-13840. doi:10.1523/JNEUROSCI.2601-12.2012.

60. de Martino B, Kumaran D, Seymour B, Dolan RJ. Frames, Biases, and Rational Decision Making in the Human Brain. 2006;684(2006). doi:10.1126/science.1127205.

61. Sharot T, Velasquez CM, Dolan RJ. Do decisions shape preference? Evidence from blind choice. *Psychol Sci*. 2010;21(9):1231-1235. doi:10.1177/0956797610379235.

62. Barraclough DJ, Conroy ML, Lee D. Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci*. 2004;7(4):404-410. doi:10.1038/nn1209.

63. Abrahamyan A, Silva LL, Dakin SC, Carandini M, Gardner JL. Adaptable history biases in human perceptual decisions. *Proc Natl Acad Sci*. 2016;113(25):E3548-E3557. doi:10.1073/pnas.1518786113.

64. Kable JW, Glimcher PW. The Neurobiology of Decision: Consensus and Controversy. *Neuron*. 2009;63(6):733-745. doi:10.1016/j.neuron.2009.09.003.

65. Romo R, Hernandez A, Zainos A. Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron*. 2004;41(1):165-173. doi:S0896627303008171 [pii].

66. Seo H, Barraclough DJ, Lee D. Lateral Intraparietal Cortex and Reinforcement Learning during a Mixed-Strategy Game. *J Neurosci*. 2009;29(22):7278-7289. doi:10.1523/JNEUROSCI.1479-09.2009.

67. Wendelken C. Meta-analysis: how does posterior parietal cortex contribute to reasoning? *Front Hum Neurosci*. 2015;8(January):1-11. doi:10.3389/fnhum.2014.01042.

68. Kiani R, Hanks TD, Shadlen MN. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci*. 2008;28(12):3017-3029. doi:10.1523/JNEUROSCI.4761-07.2008.

69. Sugrue LP, Corrado GS, Newsome WT. Matching Behavior and the representaiton of value in the parietal cortex. *Science (80-)*. 2004;304(October):457-461. doi:10.1126/science.1094765.

70. Dorris MC, Glimcher PW. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron*. 2004;44(2):365-378. doi:10.1016/j.neuron.2004.09.009.

71. Hwang EJ, Andersen R a. Effects of visual stimulation on LFPs, spikes, and LFP-spike relations in PRR. *J Neurophysiol*. 2011;105(4):1850-1860. doi:10.1152/jn.00802.2010.

72. Raposo D, Kaufman MT, Churchland AK. A category-free neural population supports evolving demands during decision-making. *Nat Neurosci*. 2014;17(12):1784-1792. doi:10.1038/nn.3865.

73. Gold JI, Law C-T, Connolly P, Bennur S. The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning. *J Neurophysiol*. 2008;100(5):2653-2668. doi:10.1152/jn.90629.2008.

74. Busse L, Ayaz A, Dhruv NT, Katzner S, Saleem AB, Scholvinck ML, Zaharia AD, Carandini M. The detection of visual contrast in the behaving mouse. *J Neurosci*. 2011;31(31):11351-11361. doi:10.1523/JNEUROSCI.6689-10.2011.

75. Hanks TD, Kopec CD, Brunton BW, Duan C a., Erlich JC, Brody CD. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*. 2015;520(7546):220-223. doi:10.1038/nature14066.

76. Hippenmeyer S, Vrieseling E, Sigrist M, Portmann T, Laengle C, Ladle DR, Arber S. A developmental switch in the response of DRG neurons to ETS transcription factor signaling. *PLoS Biol*. 2005;3(5):0878-0890. doi:10.1371/journal.pbio.0030159.

77. Komiyama T, Sato TR, O'Connor DH, Zhang Y-X, Huber D, Hooks BM, Gabitto M, Svoboda K, T K, TR S, DH O, YX Z, D H, BM H, M G, K S. Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature*. 2010;464(7292):1182-1186. doi:10.1038/nature08897.

78. Tennant KA, Adkins DL, Donlan NA, Asay AL, Thomas N, Kleim JA, Jones TA. The organization of the forelimb representation of the C57BL/6 mouse motor cortex as defined by intracortical microstimulation and cytoarchitecture. *Cereb Cortex*. 2011;21(4):865-876. doi:10.1093/cercor/bhq159.