# Representations in Rodent Primary Visual Cortex During Visual Discrimination

*A dissertation presented by*

Anqi Zhang

*to the*

Watson School of Biological Sciences

*in partial fulfillment of the requirements for the degree of*

Doctor of Philosophy

*in*

Biological Sciences

*at*

Cold Spring Harbor Laboratory

July 7, 2020

# Declaration of Authorship

I, Anqi Zhang, declare that this thesis titled, "Representations in Rodent Primary Visual Cortex During Visual Discrimination" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at Cold Spring Harbor Laboratory.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at Cold Spring Harbor Laboratory or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

_____

Date:

_____

"*The relation between what we see and what we know is never settled.*"

John Berger, *Ways of Seeing*

# *Abstract*

Although the hierarchical pathways underlying visual decision making in primates have been studied extensively, the neural circuits underlying visual decision making in rodents are not as well understood. Meanwhile, there is increasing evidence to suggest that while sensory cortices have been largely studied for their feedforward processing role, massive feedback projections distribute representations of nonsensory variables such as movement and choice across cortex, even dominating activity within early sensory cortices. The role of these representations is not well understood within a sensory discrimination paradigm. We therefore set out to investigate the representations and contributions of rodent primary visual cortex (V1) in a visual decision making task.

In this thesis, I describe a novel visual discrimination task for freely moving rats. Subjects are presented with a set of distributed flickering dots (a "cloud of dots"), and are asked to judge whether there are more dots in the upper or lower visual hemifield. Subjects report their choice with a nosepoke into the corresponding water delivery port. To facilitate experimental access in this task, I developed a software-based approach to control viewing angle at the center port by reinforcing head position using closed-loop online video tracking.

While subjects are able to find and stably carry out the comparison rule when it is necessitated by the statistics of the stimulus distribution over trials, behavioral experiments revealed that at baseline, subjects reliably converge on an abbreviated strategy, such that only half of the full stimulus is necessary and sufficient to drive the behavior. I used tetrode recordings to interrogate neuronal responses in primary visual cortex (V1) of behaving animals. In addition to classically responsive visual neurons, I report overlapping subpopulations of V1 single neurons that are not only responsive, but also selective to choice side and outcome. I show that despite a lack of clear structure to the organization of sensory and nonsensory information in V1 single neurons, the trial-to-trial variability of stimulus-driven activity depends on the behavioral relevance of visual information streams.

I further quantified the intrinsic behavioral bias towards the lower visual field. I found that animals could adapt their strategy to the statistics of their stimulus environment over many trials, but at baseline overwhelmingly preferred to base their decisions on the lower visual field, even when this strategy was suboptimal. Preliminary optogenetic inactivation experiments showed that stimulus period V1 activity was necessary for accurate discrimination behavior, while outcome period V1 activity had little effect on next trial choice. Finally, motivated by the role of corticostriatal projections in the analogous auditory "cloud of tones" task, I explored the organization of the corticostriatal projection from V1. Here I observed an asymmetry in the pattern of projections from V1 to the striatum, coinciding with V1 mapping of visual space and the asymmetry in stimulus usage by our subjects.

# *Acknowledgements*

This work could not have been completed without the support and generosity I have been lucky to receive over my years here at CSHL.

First, I would like to thank my thesis advisor, Anthony Zador, for the freedom to forge my own path in the lab. Tony consistently encouraged me to pursue my ideas, never faulted me when they turned out to be mistakes, and always credited me (often beyond what I deserved) when they led to success. Tony demonstrates daily what it means to be a rigorous thinker and a generous mentor. I could always count on Tony's confidence in me, even – especially – when I couldn't count on my own.

I would like to thank my committee members: Steve Shea, Adam Kepecs, and Bo Li. Thank you for your thoughtful input through the years. Your encouragement and intellectual engagement consistently improved both my spirits and my scientific perspective. Thank you especially to Adam, for inviting me into your lab meetings and scientific discussions.

To the staff of the Watson School of Biological Sciences – Alex, Alyson, Carrie, Monn, Kim, and Kim: thank you for tolerating my sometimes daily visits (some might call them interruptions), and thank you for the invaluable advice you have each given me. Most importantly, thank you for the possibly hundreds of dollars' worth of candy I have consumed over the years.

This work could absolutely not have been done without the support of all of the staff at CSHL. I would like to thank the LAR staff, particularly Alex, Steve, Rachel, and Kristin, for animal care and support. Thank you to Rob Eifert for building various equipment parts to my often vague specifications, on often compressed timelines.

Thank you to the entirety of the Marks building, but in particular Barry, Wiktor, and Torben, who shared invaluable knowledge and resources with me over my time here.

Thank you to all of my friends, the ones I made here, and the ones that stuck with me even as they watched me descend into the world of rodent behavior. Thank you, in particular, to Georgi, Priyanka,

and Florin, who have given me so much support, both personal and intellectual, and to Queen, Sheehan, and Evan, for offering a world beyond science.

Finally, above all, I want to thank my family for their unwavering support. Mom and Dad, thank you for your infinite love and generosity. Thank you for building a foundation for me to stand on, and a loving home to come back to. To Paul: you know how important you have been to me on this journey, but I will say it anyway. You tirelessly pick me up each time I feel like staying down. Thank you for your kindness, your love, and above all, your strength.

# Contents

# List of Figures

# List of Abbreviations

| | |
|---|---|
| A | anterior visual area |
| **AL** | anterolateral visual area |
| **AM** | anteromedial visual area |
| **AP** | anteroposterior |
| **CD** | caudate nucleus |
| **ChR2** | channelrhodopsin-2 |
| **DDM** | drift diffusion model |
| **dLGN** | dorsal lateral geniculate nucleus |
| **EIB** | electrode interface board |
| **FEF** | frontal eye fields |
| **ISI** | interspike interval |
| **IT** | inferotemporal area |
| **LI** | laterointermediate visual area |
| **LIP** | lateral intraparietal area |
| **LM** | lateromedial visual area |
| **MT** | middle temporal area |
| **P** | posterior visual area |
| **PM** | posteromedial visual area |

**POR** postrhinal visual area

**PPC** posterior parietal cortex

**PSTH** peri-stimulus time histogram

**PV** parvalbumin

**RDM** random dot motion

**RGC** retinal ganglion cell

**RL** rostrolateral

**ROC** Receiver Operating Characteristic

**SC** superior colliculus

**V1** primary visual cortex

**V2** secondary visual cortex

# Chapter 1

# Introduction

In order to survive in its environment, an animal must use its brain to extract relevant information from complex and noisy sensory scenes, and then use that information to guide appropriate responses. How animals transform sensory information into appropriate actions has been a longstanding question in neuroscience. When presented with noisy momentary evidence, sensory areas are believed to be involved in detecting stimulus features, while higher-order "cognitive" areas may integrate information over time or with other parameters such as context, reward probability, or independent information streams, to ultimately generate a decision. How sensory areas implement feature selectivity and how downstream areas accumulate information have both been extensively studied.

What remains less clear is how, and how much of, moment-to-moment activity in early sensory areas is involved in the process of decision-making. This question has become more intriguing as recent work increasingly reports that non-sensory signals, such as those related to motor activity or outcome, are both independently present and modulate sensory-related activity in primary sensory areas, particularly in rodents (Shuler and Bear, 2006, Niell and Stryker, 2010, Keller, Bonhoeffer, and Hübener, 2012, Stringer et al., 2019, Musall et al., 2019, Steinmetz et al., 2019). In some experimental paradigms, these non-sensory representations are integrated with feedforward processing to support learning and execution of sensory-guided expectations (Shuler and Bear, 2006, Keller, Bonhoeffer, and Hübener, 2012, Gavornik and Bear, 2014). The presence of expectation-modulated representations in sensory cortex suggests a possible substrate for a "cognitive filter" (Bengio, 2017), which has classically been

thought to exist only in higher cognitive brain regions, many processing stages past the sensory periphery (Reynolds and Chelazzi, 2004, Maunsell, 2015). The organization of non-sensory features relative to sensory tuning remains unknown. The structure of this organization could have implications for the ease and flexibility of implementing a cognitive filter during a sensory-guided task, such as when parsing a redundant sensory scene to facilitate efficient decision-making.

Therefore, we developed a visual discrimination paradigm to understand the organization and role of visual and nonvisual representations in V1. We wanted to understand whether this organization correlated with either innate behavioral biases toward different parts of the visual field or externally imposed task requirements. This line of questioning has some overlap with the field of active sensing, which studies changes in directed sensory acquisition based on expectation, for example, the discovery by Yarbus that human eye movements are modulated by the type of information they asked to gather about a scene (Yarbus, 1967). Additionally, the behavioral phenomenon of selectively weighting one stream of sensory information over another overlaps with the classical view of attention. We will attempt to carefully delineate our behavioral results from these other lines of questioning.

The goal of this thesis is to add to the body of knowledge about the neural representations that allow visual information to guide decisions. The experiments described in this thesis attempt to dissect the processing of visual information in a discrimination task from a behavioral, neural, and causal standpoint. This chapter will review existing work on visual processing in primates and rodents. I will also review insights from perceptual decision-making studies in primates and rodents, and the contributions of single neurons, particularly sensory neurons, to these behaviors. In Chapter 2, I describe a novel visual discrimination task for freely moving rats, as well as insights into innate asymmetries in rodent visual behavior. In Chapter 3, I further characterize these behavioral asymmetries, and show that animals can integrate information about their environment over time to adapt or override their innate behavioral tendencies in order to achieve optimality. In Chapter 4, I report single neuron data from behaving animals in both a discrimination task and a simplified behavior wherein choice is stimulus-independent. In Chapter 5, I follow up on our behavioral and neural observations with optogenetic manipulations to probe the ways in

which primary visual cortex activity drives behavior. In Chapter 6, I discuss the implications of this set of studies, and how they might be extended to achieve greater understanding of the role of early sensory cortex in behavior.

## 1.1 Visual processing pathways in primates and rodents

Visual information enters the mammalian brain through the retina. Light of different wavelengths from the visual scene is sensed via a series of molecular transformations within photoreceptor cells, which produce graded changes in their membrane potential. The resulting hyperpolarization signals through bipolar cells, with the additional influence of horizontal cells and amacrine cells, to activate retinal ganglion cells (RGCs), the first spiking node in the visual processing pathway. RGCs form the output layer of the retina, sending many parallel streams of information about features of the visual scene, such as luminance for a given wavelength, contrast, and motion, primarily to the dorsal lateral geniculate nucleus of the thalamus (dLGN) and the superior colliculus in both primates and rodents (Trenholm and Krishnaswamy, 2020). Processing in the LGN pools signals across RGCs, and prominent feedback projections suggest the integration of external signals to generate filters (Wimmer et al., 2015). LGN representations are organized by location in visual space, with their projections precisely aligned so as to produce the feature selectivity found in V1 (Reid and Alonso, 1995).

V1 neurons again pool across on-off receptive fields from the LGN to produce receptive field structures such as edge detectors. The classical oriented receptive fields first identified by Hubel and Wiesel in cat V1 are known as simple cells (Figure 1.1A), with on and off subregions that are distinct and can be linearly summated (Hubel and Wiesel, 1959). This is in contrast to complex cells, which are roughly equal in number to simple cells in primates, carnivores, and rodents, and which have more complex spatiotemporal receptive fields (Niell and Stryker, 2008, Carandini et al., 2005). V1 maintains retinotopic mapping in its representation of visual space, but in the binocular region includes information intermingled from both eyes at the level of the population. The above description of processing in the early visual

pathway is highly simplified; nonlinearities exist at all stages, giving rise to responses that are difficult to predict precisely either in degree or in kind from the visual stimulus alone.

It is well accepted that, at least in primates, two separate but interacting "streams" of processing branch from V1 to serve distinct computational goals. One of these, the dorsal stream, so named for the positions of the brain regions thought to be involved, is thought to be concerned with processing spatial relationships, while the ventral stream is concerned with object representation (Mishkin, Ungerleider, and Macko, 1983; Haxby et al., 1991). The broad categorizations arise from decades of studying specific lesion deficits, and have been backed up by recordings of cells in these areas that find distinct tuning properties across the two streams. The canonical primate visual processing streams are schematized in Figure 1.1B.

Therefore, the classical view of the visual processing hierarchy is that visual representations are transformed, by pooling and scaling responses across neurons at each subsequent stage, from amount and wavelength of light at a given pixel of visual space (retina) to detection of salient features such as edges (V1) to direction of motion (middle temporal cortex, MT) and composite objects (inferotemporal cortex, IT). In recent years, these findings have been supported by in silico work applying sparse coding principles (Olshausen and Field, 1996), and convolutional neural networks (Yamins and DiCarlo, 2016) to natural scene datasets, which produce receptive field properties and hierarchical pooling principles similar to those found in biology.

Meanwhile, feedback projections between visual cortical regions complicate the feedforward hierarchy. While careful anatomical studies of the primate visual system leveraging the laminar patterns of terminations across visual areas support a roughly hierarchical organization, they also show extensive top-down feedback projections (Felleman and Van Essen, 1991, Figure 1.1C). Within the emerging field of rodent vision, feedback into primary visual cortex has been identified not only from higher visual cortical areas, but also from non-sensory cortical areas altogether (Leinweber et al., 2017). In the rodent visual system, where the distinctness and roles of individual visual cortical regions remain less clearly defined, a well-understood processing hierarchy is still lacking.

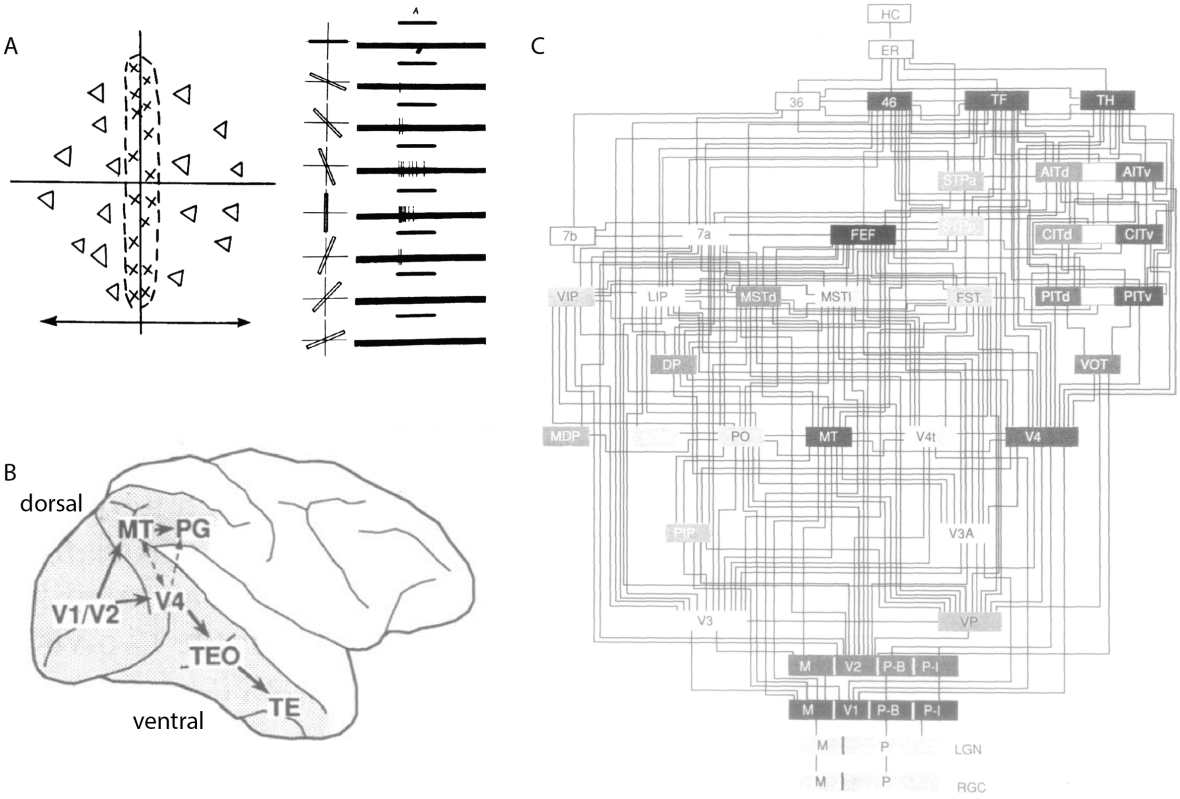


FIGURE 1.1: **Features of visual processing in primates:** A. Example simple cell receptive field mapped by presentation of dot and bar stimuli. Left: Receptive field structure. Crosses denote light positions that elicited increased activity in the recorded cell, while triangles denote positions eliciting decreased activity. Right: Responses of the same cell to different orientations of bar stimuli. Adapted from Hubel and Wiesel, 1959. B. Schematic showing the two hierarchical processing streams proposed to underlie object (ventral stream) and spatial (dorsal stream) visual processing in the primate brain. Adapted from Haxby et al., 1991. C. Connectivity map between 32 visual cortical areas and several nonvisual areas in the primate, showing hierarchical projections (from bottom to top), with most connections known to be reciprocal. Adapted from Felleman and Van Essen, 1991.

The study of visual perception in rodents is relatively new. Rodents are an appealing model system due to the greater range of genetic tools available and relatively quick breeding time. Despite having larger receptive fields than primates and lacking a high-resolution fovea, leading to much lower visual acuity, rodents have been shown to perform many visually-guided tasks previously thought to be outside their capabilities (Raposo et al., 2012, Petruno, Clark, and Reinagel, 2013, Burgess et al., 2017, Resulaj et al., 2018). Rodent work also affords the ability to probe visual behaviors and their neuronal correlates in freely moving animals, in a more naturalistic state than in classical monkey head fixation. Even head-fixed preparations in rodents permit locomotion and therefore the potential for engagement with the visual environment, for example through virtual reality (Fiser et al., 2016, Saleem et al., 2018).

In comparison to the primate, the relationships between visual areas in the rodent are not well defined. V1, located at the posterior end of dorsal cortex, is the first cortical stage of processing, receiving inputs from the LGN and containing orientation selective cells organized retinotopically, similar to primate V1, but then sends projections to three "secondary" visual cortical areas, lateromedial (LM), anterolateral (AL), and posteromedial (PM) areas (Glickfeld et al., 2013). There are indications that these secondary visual areas may serve as entrypoints to two distinct processing streams, analogous to the dorsal and ventral streams in primates, on the basis of their spatial and temporal frequency selectivity, and the targeted patterns of projections they receive from V1 (Glickfeld et al., 2013, Han et al., 2018). Additional visual areas have been identified on the basis of preserving part or all of the retinotopic map, and include areas laterointermediate (LI,) anteromedial (AM), rostrolateral (RL), postrhinal (POR), posterior (P), and anterior (A) (Wang and Burkhalter, 2007). Yet more higher visual areas have been identified in later studies (Glickfeld and Olsen, 2017). The nature of representations and processing rules in and between these areas is not yet well defined.

Rodent V1 neurons maintain many of the same feature representations found in primates, but with a few key differences. Both simple and complex cells have been found in rodent V1 (Girman, 1985). Single neurons show selectivity for orientation and spatial frequency, and maintain properties such as spatial summation and contrast-invariant tuning (Girman, 1985, Niell and Stryker, 2008). However,

receptive field size ranges typically between 5-15 degrees, approximately an order of magnitude larger than in primates. Layer 5 neurons have the poorest orientation selectivity and the largest receptive fields (Niell and Stryker, 2008). Rodent V1 contains orientation selective cells, but they are not arranged in the pinwheel organization found in primates and carnivores (Ohki et al., 2005). However, several lines of evidence have found that strength of connections between excitatory cells in V1 local circuitry (Cossell et al., 2015), as well as projections to higher visual areas (Glickfeld et al., 2013), are driven by receptive field preferences. This suggests that despite the absence of a clear physical organization, single neurons in rodent visual cortex preferentially communicate with neurons with similar tuning.

Several recent studies have found strong and reliable responses across rodent cortex, including V1, aligned to locomotion or other motor actions (Stringer et al., 2019, Musall et al., 2019, Steinmetz et al., 2019). Indeed, the magnitude of these motor responses often exceeds that of sensory responses. Similarly, modulations of visual cortical activity have also been found to correlate with estimates of arousal, derived from pupil diameter (Neske et al., 2019), although isolated periods of heightened arousal have also been reported to suppress spontaneous firing rates, presumably in service of improved visual processing by improving signal to noise (Vinck et al., 2015). How are such non-visual signals integrated into visual processing in V1? In the next section, I will describe proposed models of how nonsensory modulation and representations may be involved in specific visual computations.

## 1.2 Computational roles for nonvisual signals in rodent V1

How should we make sense of the robust motor-related activity in V1? It is believed to arise from motor collaterals that are active during self-generated motion (Leinweber et al., 2017), but a general computational role is not yet known. One hypothesis is that these signals serve an important role within a predictive model framework (Wolpert, Ghahramani, and Jordan, 1995, Rao and Ballard, 1999, Keller and Mrsic-Flogel, 2018), where motor information allows the production of sensory predictions that are then compared with sensory reality. Indeed, it has been found that activity in layer 2/3 of V1 is greater during locomotion paired with moving stimuli (drifting gratings) than during gratings presented when

the animal is stationary, but greater still when the locomotion and stimuli are decoupled from each other (Keller, Bonhoeffer, and Hübener, 2012, Figure 1.2A). Keller and colleagues term this increased activity mismatch detection, and find that these signals are driven by top-down projections from frontal areas such as secondary motor cortex (Leinweber et al., 2017). Interestingly, these topographically organized top down signals can be remapped to allow mismatch detection in a left-right inverted virtual reality, demonstrating that they are flexible and experience-dependent. In fact, the ability to generate mismatch responses depends on visuomotor coupling experience in early development (Attinger, Wang, and Keller Correspondence, 2017). Further experiments revealed that mismatch detection is built up across visual space similarly to the visual scene: using single neuron receptive fields that signal local deviations from predicted visual flow (Zmarz and Keller, 2016). In addition to supporting locomotion-paired visual flow, these layer 2/3 mismatch neurons also respond robustly to the omission of an expected stimulus at a trained location in virtual reality (Fiser et al., 2016).

Locomotion is closely related to spatial location, especially in a reward-guided navigation task. The ability of individual cells to signal the unexpected absence of a stimulus in physical space suggests V1 may contain independent information about spatial location. Indeed, a recent study showed that non-visual task variables such as subjective spatial location can also differentiate responses to identical visual stimuli during virtual navigation (Saleem et al., 2018). These modulations of stimulus-driven responses are independent of motor differences, such as running speed, and produce at the population level an estimate of the animal's current subjective position, correlating with behavioral measures. The encoding of subjective position may partially derive from or occur in tandem with encoding of 3-dimensional head orienting movements, which have been reported to modulate V1 activity in a direction-specific manner (Guitchounts et al., 2020).

Similar to other modalities (Jaramillo and Zador, 2011, Xin et al., 2019), V1 representations have been shown to adapt to behavioral demands over the course of learning, correlating with improved discrimination of task-relevant stimuli (Poort et al., 2015). In this study, discrimination of visual patterns (oriented gratings) improved at the level of the population over training on a virtual navigation task, in

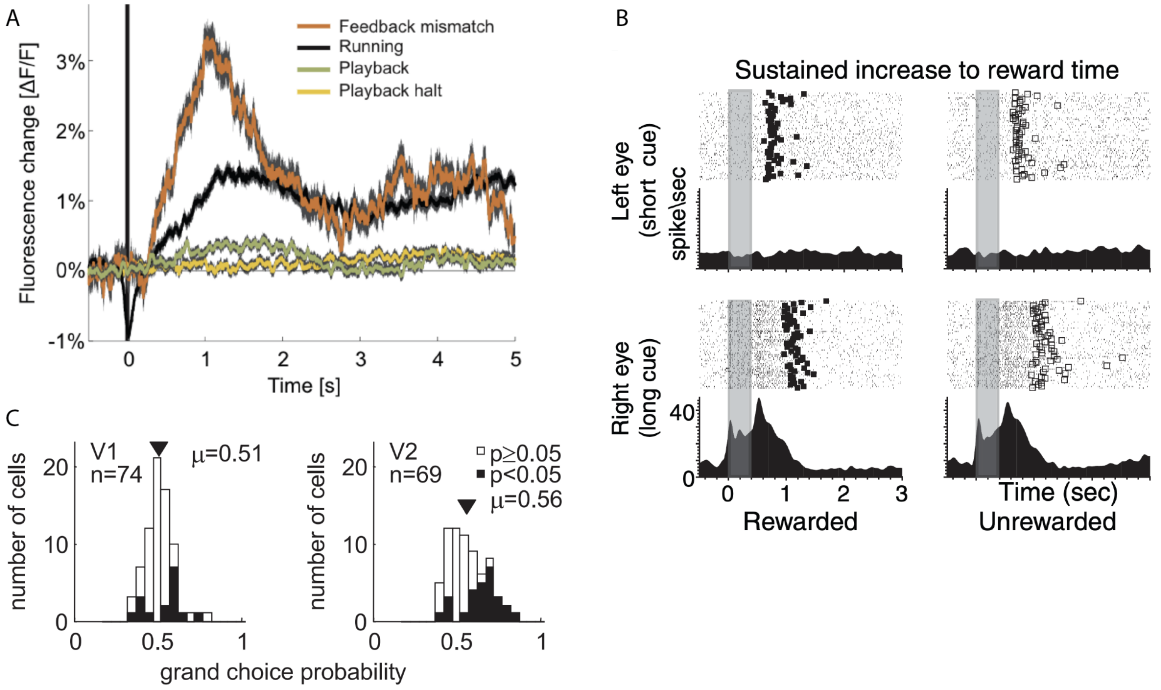


FIGURE 1.2: **Nonsensory signals in V1:** A. Modulation of population bulk visual response (playback, green trace) magnitude by locomotion (running, black trace) and expectation mismatch (feedback mismatch, orange trace). Adapted from Keller, Bonhoeffer, and Hübener, 2012. B. Representation of expected reward timing in single neuron firing. This example neuron shows sustained firing until time of expected reward, including on unrewarded trials, specific to one of two stimulus conditions. Exact timing of reward was determined by the number of licks received. Adapted from Shuler and Bear, 2006. C. Significant choice probabilities during a disparity discrimination task (judging near versus far) using random dot stereograms are fewer in V1 than in V2 (shaded bars). Adapted from Nienborg and Cumming, 2006.

a task-engagement dependent manner, by recruiting previously non-selective cells to become selective for the rewarded orientation. Additionally, over training, choice and anticipation signals appeared in a subset of neurons that preferred the rewarded stimulus.

Potentiation of responses to learned visual patterns has also been shown using temporal sequences of stimuli. Gavornik and Bear (2014) showed that trained sequences of oriented gratings elicit greater responses in layer 4 of V1 than novel sequences, and furthermore, that these responses are attenuated when trained sequences are manipulated in serial order or timing. Similarly, when single elements of the sequence were omitted, cortical responses were smaller but still significantly potentiated relative to a novel sequence. This result appears at first to contradict the predictive algorithm found in the locomotion domain; a "mismatch" from the trained sequences produced a lesser, rather than greater, response. However, in this case, the stimulus sequence is not tied in closed-loop to the animal's own movements, suggesting that predictive signals in V1 differ in sign, and therefore possibly computational role, depending on the precise behavioral demands and relationship to the visual scene.

Response magnitude is not been the only aspect of V1 activity that is modulated by training or behavioral relevance. In a visually-cued timing task, single neurons preferring a given stimulus sustained, ramped, or inhibited their firing until the time of expected reward associated with that cue, even on catch trials when reward was withheld (Shuler and Bear, 2006, Figure 1.2B). Further work on a similar task where the timing of the first lick determined the reward magnitude showed that the variability in V1 responses held information about the trial-by-trial waiting time, and that this behavioral timing could be shifted by optogenetic perturbation of V1 neuronal activity (Namboodiri et al., 2015).

A number of studies have attempted to distinguish pure sensory processing from task-related sensory signals, by comparing responses within subjects across passive and active sensation tasks. A recent study that recorded brain-wide single neuron spiking activity in head-fixed mice found that cortical baseline activity, including in V1, was decreased during task engagement compared to passive viewing (Steinmetz et al., 2019), while stimulus-evoked activity in V1 remained similar independent of engagement. In the auditory domain, it has been shown that frequency selectivity shifts from uniform encoding of the

stimulus frequency across the population during passive listening, to increased selectivity to stimuli near a trained category boundary after learning and performance on a task (Xin et al., 2019). It is important to note that these comparisons were made in head fixed animals that are in engaged versus disengaged states within their environment. However, finer behavioral differences between these states are difficult to quantify. While experimenters typically consider and try to regress out the role of arousal, often measured by pupil diameter, locomotion, or other movement patterns, there can remain other fine behavioral or motivational changes that occur systematically between these two states that are not easily quantifiable by video tracking.

Together, this body of work has made significant inroads in uncovering how representations in primary visual cortex both reflect and support behavioral demands. However, much remains to be understood about the generality of these proposed computations. The relationship between encoding of task-related variables and sensory encoding has not been well studied at the single cell level, though one study has suggested that choice signals are gradually enriched depending on stimulus preference over the course of learning (Poort et al., 2015). Additionally, the role of these non-sensory representations in the context of a sensory-guided decision task remains unclear. It remains the case that much of our understanding of how decisions based on sensory evidence are generated is based on the primate literature. In the next section, I will review main points from the primate literature, and discuss additional insights from more recent rodent circuit work during perceptual decisions.

## 1.3 Neural substrates for perceptual decision-making

The limits of information processing in early sensory systems can be inferred through physiology, through measurements of variables such as the size of a receptive field or the latency of an event-triggered action potential. However, while these measurements provide an upper bound on the information available to the brain, they do not directly translate to the limits of cognitive processes such as perception, valuation, or learning. Behavioral reports are therefore required to understand what kinds of computations brains implement and how they do so.

Perceptual decision-making probes the mechanisms of perception by asking subjects to report judgments based on the perceptual information available at a given instance. These decisions, made with different amounts of information, and the neuronal correlates of this process in perceptual and associative areas, provide clues about how perceptual information is integrated, categorized, and associated with an appropriate motor output. Extensive work on the now-canonical Random Dot Motion task in primates has identified areas that represent the sensory evidence, the amount of evidence currently available, and the likely choice on a given trial (Britten et al., 1992, Roitman and Shadlen, 2002, Gold and Shadlen, 2007). In this task, head fixed monkeys were presented with a collection of moving dots, of which a certain percentage moved in a coherent direction, while the remaining dots moved in random directions. The monkey was rewarded for judging the correct direction of the coherent dots, which it reported with a ballistic eye movement (called a saccade) to one of two target sites. In this fixed time task, early work identified single MT sensory neurons that matched or exceeded the behavioral accuracy of the animal, and further, found that the trial-to-trial variability of single neurons in MT weakly predicted trial-to-trial behavioral variability (Britten et al., 1992, Britten et al., 1996). However, later studies using more constrained estimates of length of stimulus use found that single MT neurons were less accurate than the population or the animal at solving the task (Cook and Maunsell, 2002).

This well-defined behavioral paradigm has produced three decades of work to attempt to understand where this sensory information is integrated over time, and how representations change as this occurs. The lateral intraparietal area (LIP) was identified as a potential site for integration of pure sensory evidence into what has been termed a "decision variable," a representation that links sensory inputs to a categorical choice which then manifests as a motor report. LIP is well-positioned to serve this role, as it receives projections from extrastriate visual cortex, and projects to oculomotor regions that would be responsible for the saccadic choice report (e.g. frontal eye fields, FEF, and superior colliculus, SC). Further, it is organized by the direction of oculomotor movements, rather than in coordinates of visual space, suggesting its representations are well-suited to convert sensory input into a movement-specific signal. Indeed, LIP single cell activity showed stereotyped ramping patterns whose slopes corresponded

to the strength of the moment-to-moment motion stimulus, and were sensitive to brief pulses of incoherent motion perturbations delivered during the stimulus period. Later work found similar activity profiles in FEF, SC, and caudate nucleus (CD), the input nucleus of the primate basal ganglia (Kim and Shadlen, 1999, Roitman and Shadlen, 2002, Ding and Gold, 2010). Such ramping cells are thought to be neural substrates for evidence accumulation over time, and the alignment of their trajectories on reaction time tasks with behavioral measures of decision formation gave support to the drift diffusion model (DDM) of evidence accumulation. Under such a model, two opposing streams of evidence are independently accumulated in selective neural populations from momentary evidence, until one of these populations' activity reaches a threshold, at which point a decision is made. Further microstimulation studies manipulating the activity of the above populations found effects on choice bias and reaction time that aligned with the expectations of the DDM given a change in the starting value of the accumulator.

Rodent decision-making work has shown that rats and mice, like primates, are capable of accumulating perceptual evidence and reporting their decisions (Sanders and Kepecs, 2012, Brunton, Botvinick, and Brody, 2013). From a body of work spanning many input modalities (audition, vision, olfaction) and many types of choice reports (lick port, wheel movement, virtual navigation), several areas have been implicated in the rodent brain as having a role in linking noisy sensory information to appropriate choices. Work from the Brody lab found that rats are able to accumulate pulses over randomly timed competing auditory trains with zero instability or leak under a drift diffusion model (Brunton, Botvinick, and Brody, 2013). They found that under this paradigm, all clicks were integrated with the same weight across the duration of the stimulus. Accumulation of sensory evidence in rodents has been found in some cases to operate at quick timescales and with high accuracy. One study found that mice are able to base decisions on integration of <100ms of visual information (Sriram et al., 2018), while another study estimates that the first 80ms, corresponding to the first few spikes in visual cortex, are sufficient to allow orientation discrimination (Resulaj et al., 2018).

Representations of the accumulated evidence have been found in population responses recorded from frontal and parietal cortices (Hanks et al., 2015; Scott et al., 2017). Due to its homology to primate

LIP, rodent posterior parietal cortex (PPC) has been hypothesized to be an important site for evidence integration during decision-making. This comparison has been bolstered by single neuron recordings in rodent PPC finding encoding of accumulated evidence, navigation trajectory, history dependence, and choice (Hanks et al., 2015, Morcos and Harvey, 2016, Brody and Hanks, 2016, Licata et al., 2017, Krumin et al., 2018, Akrami et al., 2018).

Subcortical nuclei have also been implicated in mediating sensorimotor transformations. The striatum has received particular attention, perhaps due to the proposed motor execution role of the basal ganglia circuitry for which it serves as an input nucleus. The striatum is a large subcortical structure that receives topographically organized inputs from essentially all cortical areas, and has been partitioned into several hypothesized functional subregions. Ventral striatum, which receives dense inputs from the dopaminergic neurons in the ventral tegmental area, is thought to underlie value-based decisions. The dorsal striatum, also known as the caudate in primates, and further divided into dorsomedial striatum (DMS) and dorsolateral striatum (DLS) in rodents, is classically thought to mediate action selection and initiation. The DMS has been proposed to support flexible goal-directed behaviors, while the DLS is thought to implement habitual reflex-like behaviors. However, finer segmentations of the striatum can be defined based on the patterns of projections from cortex (Hunnicutt et al., 2016) and the resulting representations in striatal neurons. For example, primary visual cortex has been shown to project to dorsomedial striatum in mice (Khibnik, Tritsch, and Sabatini, 2014), while primary auditory cortex projects to posterior striatum in rats and mice (Xiong, Znamenskiy, and Zador, 2015). Functionally, DMS neurons have been shown to be responsive to visual stimuli, while whisker deflection responses are detected in the DLS (Reig and Silberberg, 2014, Sippy et al., 2015).

The behavioral role of these striatal subregions may depend on the pattern of cortical inputs they receive. Znamenskiy and Zador (2013) showed that axons in the posterior striatum, which receives inputs from auditory cortex, can bias behavior towards a choice associated to the preferred stimulus of the subset of corticostriatal neurons stimulated with channelrhodopsin (ChR2, Znamenskiy and Zador, 2013). Distinct from earlier primate work (Ding and Gold, 2012), the bias direction here was correlated with the

local sensory tuning of the input to the striatum, rather than anti-correlated with the choice preference of the striatal neurons themselves. Further work (Xiong, Znamenskiy, and Zador, 2015) showed that auditory projections to this region of the striatum are organized by tone frequency, and the synaptic strength of these projections across this tonotopic map forms a gradient in trained animals that coheres with the sensorimotor contingency the animal has been trained on (high frequency – left choice, for example). This work suggests a possible sensorimotor transformation circuit that does not require integration over time by associational cortices, but instead where choice behavior is linked directly to representations in early sensory cortex. It remains an open question whether this corticostriatal circuit motif supports only auditory discrimination behaviors, or is general across sensory modalities; additionally, it is unknown whether the structure of the sensory stimulus is important in recruiting this processing circuitry, and whether this circuit logic applies across species.

Modelling studies have proposed distinct computational roles for subcortical projection pathways. In one spiking neuron model, the effects seen by adjusting the synaptic strengths of either cortico-collicular projections or corticostriatal projections are consistent with a hypothesis where the cortico-collicular pathway detects threshold crossing by the activity of the accumulator population, while the corticostriatal pathway sets the decision threshold (Lo and Wang, 2006). This hypothesis may begin to reconcile the seeming disparity between the need for an accumulator population in the DDM, versus the apparent readout of stimulus-choice association through synaptic strength in the corticostriatal pathway following training on the above auditory discrimination task (Xiong, Znamenskiy, and Zador, 2015).

## 1.4 Single sensory neuron representations during decision-making

The single neuron is frequently thought of as the basic unit of brain computation and function (Parker and Newsome, 1998). The classical study of receptive fields suggests that the computational roles of single neurons correspond to the features that drive their activity, or what they "represent". This representational point of view has influenced the study of brain function extensively, and has especially driven our understanding of functions and computations at the sensory and motor peripheries. The study

of representations has also been applied to more abstract features, including derived and intermediate computational quantities including spatial location, decision variables, reward value, reward probability, and metacognitive quantities such as confidence (O'Keefe, 1976, Roitman and Shadlen, 2002, Gottfried, O'Doherty, and Dolan, 2003, Samejima et al., 2005, Ott, Masset, and Kepecs, 2018). Though these representations and their associated functions typically are most prominent in nonsensory brain regions, including parts of parietal and frontal cortices as well as subcortical nuclei, such features have also been found to modulate responses in sensory cortices. Below, I will review the literature on representations in sensory neurons during decision making.

The contribution of single neurons in sensory cortex to decision-making has been primarily studied through the lens of their representations of the sensory scene. While early work in MT suggested that single neurons could be either less or more accurate in their discrimination of stimuli than the full animal (Britten et al., 1992), later work limiting the stimulus delivery time (and therefore the time period over which neural activity was averaged), found that all neurons in MT were weaker in their discrimination than the animal's behavior (Cook and Maunsell, 2002). This suggests that selective pooling and optimal decoding may be required to read out the sensory information in downstream accumulation and choice areas. Similar approaches have been taken to understand how encoding of somatosensory stimuli contributes to stimulus frequency discrimination (Romo and Salinas, 2003), as well as in other visually-guided tasks such as disparity discrimination (Shiozaki et al., 2012).

A related question is whether all neurons that respond to the currently relevant stimuli contribute to the decision process. One measure that reflects a correlation between a neuron's firing and the choice behavior of the animal is the choice probability. The choice probability measures the extent to which the variability in stimulus-evoked activity in a single neuron can predict the upcoming decision. This is especially informative for ambiguous stimuli, where resulting choice reports are closer to chance levels. It has been found that choice probabilities increase as one moves through the processing hierarchy. At the sensory periphery (e.g. V1), few neurons show significant correlation in their firing with the eventual choice of the agent, while in higher sensory areas (V2: Nienborg and Cumming, 2006, V4: Shiozaki

et al., 2012, MT: Britten et al., 1996, MST: Celebrini and Newsome, 1994), more neurons display significant choice probabilities (Figure 1.2C). However, the presence of significant choice probabilities does not confirm that a given neuron is causal for a behavioral choice. Choice probabilities appear to arise from choice feedback signals from higher cortical areas (Roelfsema and Lange, 2016, Nienborg and Cumming, 2009, Yang et al., 2015); while readout of this feedback has been shown to be correlated with stimulus sensitivity in the somatosensory system (Yang et al., 2015), there is no direct evidence that these choice signals specifically target causal neuronal populations or influence their processing.

Direct manipulation of neuronal activity during behavior has become steadily more feasible over the past several decades. Early work in the RDM task found that electrical microstimulation of a small number of neurons in MT led to increased detection of the direction towards which those neurons were tuned, relative to baseline behavior, supporting a direct causal relationship between motion direction tuned cells and the eventual choice report (Salzman et al., 1992). More recently, optogenetic techniques, which harness the power of excitatory and inhibitory light-gated channels to increase or decrease the activity of affected cells, have made it possible to conduct these experiment at the broader population level. A recent study that unilaterally inactivated 9 sites in visual cortex found a similar degree of bias against the contralateral choice across all sites within visual cortex for a given hemisphere (Zatka-Haas et al., 2019). Inactivation or perturbation of V1 activity has been shown to affect performance on detection tasks, different types of discrimination tasks, and visually-guided reward timing tasks (Petruno, Clark, and Reinagel, 2013, Namboodiri et al., 2015). However, these studies did not directly show dependence of behavior on visual tuning properties at the single cell or local population level.

Though much is known about the response properties of sensory neurons during stimulus presentation, sensory neurons also show activity during epochs when no stimulus is presented. Steinmetz and colleagues find that single neurons throughout the brain, including in early sensory cortices such as V1, reliably respond immediately preceding choice initiation. However, they report that tuned choice-selection responses leading up to choice initiation are not present in sensory cortices, but only in frontal

and subcortical motor regions, such as striatum and M2 (Steinmetz et al., 2019). The behavioral relevance of choice initiation signals across cortex, including in sensory cortex, is not known.

## 1.5 Stimulus selection during behavior – attentional and other views

Sensory-guided behaviors in the natural world require an animal to differentially weight incoming streams of information, on the basis of their reliability, salience, or the animal's own current motivational state. By design, most laboratory behavioral paradigms feature a highly reduced sensory environment that contains only immediately relevant information. While stimulus use is not uniform even in these tasks, for example as seen in the fact that many decisions are predominantly based on early sensory evidence (Morcos and Harvey, 2016), the most systematic way of studying the effects of stimulus selection by attentional mechanisms has been through cued spatial attention tasks.

Cued spatial attention tasks are typically built upon a visual detection or discrimination task, such as asking the subject to detect a slight change in orientation or the presence of a low contrast stimulus. The behavior on this non-attentional version of the task is then compared to behavior following attentional cueing to a region in visual space, typically indicated by a localized luminance or color change in the cued region during a separate cueing epoch prior to stimulus delivery. In humans and non-human primates, both the attentional cue and stimulus are delivered while subjects fixate on a central point, such that behavioral enhancement is not due to increased foveation of the attended location. Attentional cueing in this paradigm has been shown to improve detection latencies and thresholds, as well as discrimination sensitivity for features such as orientation (Reynolds and Chelazzi, 2004).

Visual processing changes underlying attentional modulation has predominantly been studied in nonhuman primates. It has been proposed that changes at the neural level effectively implement a top-down filter on the sensory system, which may be useful for directing limited processing resources to the most informative features or spatial location. The evidence that this filter is implemented in a top-down fashion from brain regions downstream of sensory areas comes from observations that responses in area

LIP correlate with increased sensitivity, while direct microstimulation in FEF increases both behavioral and neural sensitivity in V4 (Reynolds and Chelazzi, 2004, Armstrong and Moore, 2007).

There are a few classical signatures observed in sensory cortices across attentional studies. When stimuli in the receptive field of a neuron are attended to following cueing, evoked responses typically increase, and their variability and pairwise noise correlations decrease (Spitzer, Desimone, and Moran, 1988, Mitchell, Sundberg, and Reynolds, 2009, Cohen and Maunsell, 2009). The proposed computational role of gain modulation and changes in noise correlations is to increase the signal-to-noise ratio of sensory representations for the attentionally primed area of visual space, and therefore improve performance based on the information in that area. Other effects observed in V4 include changes in the dynamics of transitions between high- and low-activity states (on-off states, Engel et al., 2016).

Meanwhile, effects related to attentional modulation can be seen as early as the LGN-V1 synapse: one study found that when spatial attention was directed toward the receptive field of a recorded layer 4 V1 neuron, the probability of driving that neuron to fire through monosynaptic LGN inputs significantly increased, suggesting modulation of synaptic efficacy by top-down attention mechanisms at this early processing stage (Briggs, Mangun, and Usrey, 2013). While there were only weak effects on firing rate across the recorded V1 population as a result of attention, the proportion of synchronous spikes in V1 increased. However, other classical signatures of attentional modulation have not been reliably observed in V1. For example, while spike count ("noise") correlations between V4 neurons decrease as behavior improves following attentional cueing, V1 noise correlations are not correlated with performance (Ni et al., 2018). Changes in the firing rate variance, measured as the variance of the spike count over the mean spike count, also called the Fano Factor, have been found consistently in V4 (Mitchell, Sundberg, and Reynolds, 2007, Cohen and Maunsell, 2009), MT (Niebergall et al., 2011), and more variably in V1 (Herrero et al., 2013, but see Hembrook-Short, Mock, and Briggs, 2017).

Though attentional modulation has been studied experimentally by separating the stimulus feature of interest from the process of selecting where attention ought to be placed, in more complex tasks, the two processes of stimulus selection and stimulus-guided decision-making may be interdependent. An

extension of the concept of top down filtering of V1 activity from a discrete process as in attentional cueing to a continuous process is the hypothesis that V1 may serve as a "cognitive blackboard" (Roelfsema and Lange, 2016). This framework postulates that V1 may allow continuous read-write access to higher visual areas, and therefore its activity can reflect intermediate computations that are fed back from downstream areas. There is some evidence for this hypothesis. In a curve tracing task, V1 responses are amplified for locations in visual space corresponding to the traced curve, even when this is the incorrect curve (Roelfsema and Lange, 2016, Roelfsema and Spekreijse, 2001). In the same task, the reward value of individual curves acted similarly to attentional cueing to modulate responses (Stanisor et al., 2013), suggesting that top-down modulation of early sensory cortex may be accessible to many systems, either through convergent or distinct mechanisms.

While covert attentional modulation is one way to selectively weight incoming information on the basis of its expected task relevance, as in many of the cases above, overt attentional mechanisms also play a large role in allowing animals to filter their sensory scene. Overt attentional mechanisms, such as directed eye movements to train the high acuity fovea on an attended object or features, can be considered a form of active sensing, which refers to directed movements guided by incoming sensory information to allow more efficient sensory sampling (Yang, Lengyel, and Wolpert, 2016). In animals with a fovea, including humans and non-human primates, eye movement tracking studies have consistently found that eye fixations spend the most time in the most informative regions of a visual scene (Yarbus, 1967, Toscani, Valsecchi, and Gegenfurtner, 2013, Peterson and Eckstein, 2012). Comparable studies have been few in non-foveated animals like rodents, where eye tracking studies show the majority of eye movements serve to compensate for head movements, and decrease substantially under head fixed conditions (Wallace et al., 2013, Meyer, O'Keefe, and Poort, 2020). Given the strong connection between head and eye movements, it is possible that head orientation in rodents may serve a role in visual sampling, such as through documented "head initiated saccades," where eye movements follow briefly in the same direction as head movements (Meyer, O'Keefe, and Poort, 2020). Active sensing is better-described for other sensory modalities in rodents, such as somatosensation and olfaction (Kleinfeld, Ahissar, and Diamond, 2006, Wachowiak, 2011).

The study of attention in rodents is still in its early stages (Wang and Krauzlis, 2018), and most of what is known about how attentional modulation affects sensory processing comes from primate work. The presence of robust non-sensory responses in rodent V1 offers an opportunity to gain a broader perspective on the effects of attentional modulation on activity patterns and information processing in primary visual cortex. For example, do the attentional signatures found in sensory responses also apply to non-sensory representations? Are non-sensory representations modulated or gated in any way by attentional mechanisms? The diverse representations in rodent cortex may provide new avenues of inquiry into mechanisms of attentional selection.

## 1.6 What we talk about when we talk about V1

Given the broad connectivity between V1 and the rest of the brain, as well as the diversity of responses found in V1, what do we mean when we talk about V1? Surely we can no longer think of V1 in terms of only its feedforward processing of oriented luminance changes across visual space. As evidence that feedback drives a large proportion or perhaps even the bulk of V1 activity, the field is moving towards a different framework. We might therefore talk about V1 from the point of view of the representations it contains. In this domain, though we understand fairly well the visual parameters encoded in primary visual cortex, we still lack a detailed understanding of the degree and specificity of its non-visual representations, as well as how these representations interact locally with visual representations. Similar to how early visual studies were driven by understanding the organization of visual space and orientation selectivity across visual cortex, it would be helpful to understand the functional organization of nonsensory representations across V1, and to what extent this is flexible to the visual goals of the animal. Given the targeted organization of feedforward circuits on the basis of visual tuning (Hubel and Wiesel, 1972, Glickfeld et al., 2013), how do nonsensory representations relate to and influence sensory representations at the single cell level? Do these response patterns depend on the behavioral state or goal of the animal? Answering such questions will begin to provide clues to the computations that are possible in V1, as well as what types of information might be made available to downstream areas.

But ultimately, to understand V1 requires not only knowing what V1 represents, but what it contributes to behavior. For this, we need to talk about V1 from the point of view of its influence on downstream areas, and the functional roles of the neurons within it. While normal V1 activity has been shown to be required for a number of tasks, we still lack evidence for a direct relationship between a neuron's visual tuning and the behavior it is expected to drive, in the way that has been shown for MT, much less the causal effect of a neuron's nonsensory activity. The combination of these two approaches to V1 may produce a more complete computational role for the first cortical stage of visual processing, than that of the classical edge detector.

## 1.7 Outline

In this thesis, I will first describe a novel visual discrimination task for freely moving rodents. I will show through a series of behavioral manipulations that animals are able to solve this task using only a subregion of the stimulus, and indeed do so in a consistent manner. I then show that, despite suboptimal reward return in a modified version of this task, animals still converge over time to overweight the preferred subregion. I will show that animals' strategies are sensitive to the distribution of stimuli seen over time. Then, I will investigate the organization of stimulus, motor, and outcome representations in V1, and whether they depend at the single neuron and population level on whether the preferred stimulus is causal (i.e. used or unused) for the animal's choice. Finally, I describe optogenetic experiments to probe the causative role of V1 in driving this learned discrimination behavior. While inactivation of V1 during the stimulus period has a significant effect on the choice behavior on those same trials, inactivation during the outcome period does not significantly affect next-trial choice.

Together, I will attempt to tackle the question: how is information, both sensory and nonsensory, encoded, organized, and used in V1 in the context of a learned discrimination behavior? Previous reports of motor-related and reward-related activity in V1 have established the presence of behaviorally driven non-visual signals in V1. Here I present a learned behavioral paradigm that can be used to extend these findings.

# Chapter 2

# "Cloud of Dots" Visual Discrimination Task

## 2.1 Introduction

In the introduction, I reviewed recent reports of various task related signals in visual cortex, including primary visual cortex. However, the functional role of these representations during a freely-moving visually-guided decision task is incompletely understood. Indeed, evidence linking even stimulus representations in V1 directly to behavior is sparse. Therefore, I wanted to probe V1 activity during a visually-guided behavior. To this end, I designed a task adapted from an existing auditory task used in the lab (Znamenskiy and Zador, 2013), maintaining the task structure but changing the sensory modality. Here, I ask animals to judge the spatial distribution of a distributed flickering dot stimulus. Below, I will describe the task design and trial structure, and describe features of choice behavior on this task.

## 2.2 Task design

To understand the nature of processing in primary visual cortex during visually guided behaviors, I designed a task to probe the visual discrimination abilities of freely moving rats. I used a three nosepoke setup to allow animals to report decisions based on the spatial distribution of a visual stimulus (Figure

2.1A). This setup allowed animals to self-initiate trials by poking into the center, stimulus delivery port. Animals were required to maintain their nosepoke through a variable delay to trigger stimulus delivery. The spatially distributed stimulus presented an opportunity to exploit the retinotopic organization in V1 for later experiments and manipulations; to enable this, the stimulus must be oriented in a reproducible manner in the animal's visual field over trials. Therefore, I also required animals to fulfill a head position criterion prior to and throughout the stimulus period (details below). A decision tone signified successful stimulus sampling, and the animal was able to initiate a choice movement by exiting the center port and moving into one of the two side ports. Following a correct side port choice, animals were given a small water reward (20 μL). In order to trigger reward, animals were required to sustain their side port nosepoke for a short duration of 50ms, to discourage ballistic sampling of the choice ports. Following an incorrect choice, animals were given a punishment signal in the form of a white noise burst and a short time out (between 5 to 8s). Trials aborted prior to the decision were given a shorter time out of 2s. Missed rewards did not trigger any punishment or time out; the next trial was immediately available to the animal.

I placed a monitor directly in front of the three ports, and raised it to a height such that the bottom edge was not occluded by the ports. I presented a stimulus of randomly distributed dots that flickered over frames. Animals were trained to report whether the lower half (henceforth called the lower hemifield, though the stimulus does not fill the full visual field) or upper half (upper hemifield) of the stimulus had a greater density of dots over the fixed stimulus duration (500ms). Animals reported their percept with a poke into the associated side port, according to the stimulus-action contingency they were trained on. For the easiest stimuli, trained animals consistently performed more than 95% correct.

The training protocol proceeded as follows. Animals were habituated to the behavior arena, and then trained on this task in stages. During habituation, animals were presented with a constant light over the center port, to compel them to explore the ports. When animals poked into the center port for any duration of time, any following side port visit led to a reward. Most animals were able to learn this operant behavior within a session or two. Following >200 successful executions of this behavior on a single session, animals were exposed to the easy stimuli during their time in the center nosepoke. Following the

center nosepoke, the animals were now only rewarded for making the correct choice associated with the stimulus that was just presented. The duration for which animals were required to remain in the center nosepoke gradually increased from 200ms to 500ms over training sessions. Additionally, the duration for which they were required to wait in the side port before a reward would be issued was increased from 30ms to 50ms. When animals acquired the stimulus-choice association to >80% accuracy, first intermediate difficulty trials, then finally difficult trials were introduced over the next sessions.

## 2.3 Virtual head fixation using video tracking

I wanted to maintain (roughly) the same viewing angle for animals across trials and sessions. To do this, I implemented a video tracking algorithm to gate trial progression dependent on the position of the animals' head at the center port (Figure 2.1B). Briefly, I used an IR camera and IR light source to view the animal's position over the course of the session. When the animal triggered trial onset by poking in at the center nosepoke, the head position detection algorithm was also triggered. I used ear position as a proxy for measuring the position and angle of the head. While the animal maintained its poke in the center nosepoke, the algorithm detected the presence and relative position of the two ears within two regions of interest. Once a distance criterion was met (e.g. y-distance between centroids is less than a threshold value), the head was considered to be in alignment, and the trial progressed as usual. If at any point during the pre-stimulus delay, stimulus period, or post-stimulus delay, either the x-distance or y-distance criterion was broken, the trial was aborted and the animal was accordingly given a 2s time out. There was no direct feedback for misalignment, apart from trial abortion and therefore the lack of a following decision tone or feedback for a side port choice.

Animals were trained on this rule in stages. The first (optional) stage of training required a single momentary alignment event during the stimulus presentation and associated delay periods. However, I observed that ability to achieve this criterion did not necessarily translate to learning the rule that the head must remain stationary during the duration of the nosepoke; therefore, I frequently skipped this step. From here, I enforced the full head fixation criterion, but for shorter stimulus durations that I

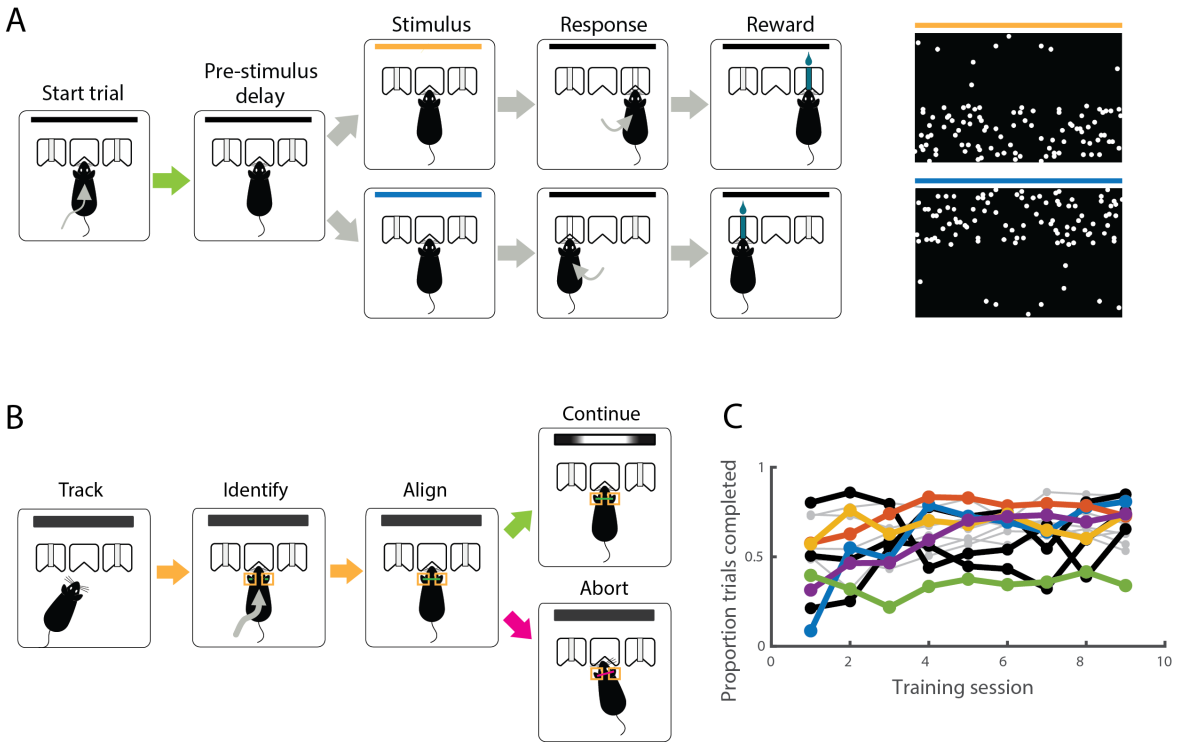


FIGURE 2.1: **"Cloud of Dots" task design:** A. Task design, with example stimulus frames for 95% lower hemifield (top) and 95% upper hemifield (bottom) trials. B. Virtual head fixation algorithm, expanded task structure from green arrow in A. C. Proportion of trials completed increased as animals were trained on head fixation. Across all animals, the mean proportion of completed trials on day 1 of training was 0.498; this increased to 0.679 by day 9. Colored lines indicate animals who were later implanted with tetrode drives in the cloud of dots task, while black lines indicate animals later implanted with tetrode drives in a simplified viewing task (Section 2.6). Grey lines indicate animals who were later implanted with fibers for optogenetic experiments (Chapter 5).

then increased over sessions. For example, the first day of head fixation training required the animal to achieve the alignment criterion throughout the pre-stimulus delay, a 200ms stimulus (rather than the full 500ms), and a post-stimulus delay. Following successful acquisition of this rule, I then extended the stimulus period gradually to its full duration, thereby also extending the duration of time between head alignment onset and the decision tone (the point at which fixation can be broken). Most animals were able to return to full stimulus sampling duration in less than a week.

Many animals were able to adjust their approach behavior to improve performance on this rule within the first 3 sessions, even as I increased the required duration (Figure 2.1C). There was heterogeneity within the population, and some animals took upwards of 10 sessions to learn to consistently hold the proper head position for the full duration period (green line in Figure 2.1C). I also observed that the magnitude of deviations from the target head position decreased over sessions. While the closed loop control over trial progression filtered out trials with disqualifying head movements, learning over exposure to this rule helped animals to increase the number of trials they completed per session.

Once animals were trained on this head positioning rule, I continued to reinforce this behavior by sustaining this requirement throughout all ensuing experiments. Though it was evident from their approach and movements that animals acquired and stably carried out this rule, the combination of this criterion with a non-zero probability of early withdrawals from the port led to a variable proportion of missed trials within any given session. For this reason, I chose to train animals on head position alignment only after they learned the three-port task structure and the sensory discrimination rule. I found no deterioration of animals' performance on the discrimination task following head position training; indeed, possibly due to increased training on the task, behavior tended to improve (higher sensitivity, less bias) in later sessions, where the head fixation protocol was active (Figure 2.2B right panel).

## 2.4 Choice behavior, but not latency, reflects stimulus difficulty

I began by training animals, after rig habituation, on the easiest stimuli available (the 5%-95% and 95%-5% distributions). The amount of time required to acquire this rule varied across animals, but all animals

were able to learn to higher than 90% accuracy between 1-4 weeks (Figure 2.2A), with a median training time of 14 sessions.

Upon satisfactory performance on the easiest trials, I began to probe behavior at 8 intermediate difficulty levels. The number of individual dots was kept constant across all trials and all stimulus frames. This meant that the upper visual field strength was always complementary to the lower visual field strength, with the total luminance across the screen kept constant across all trials.

Animals' choices reflected the stimulus difficulty, producing psychometric behavior within and over sessions (Figure 2.2B). However, contrary to what has been reported in reaction time discrimination tasks, I did not observe any significant increase in "reaction time" (either latency to initiate movement, Figure 2.2C, or speed of movement between ports, Figure 2.2D) for difficult versus easy trials. Individual animals showed slight idiosyncratic differences in their approach speed to each of the two choice ports, but there appeared to be no systematic contribution of trial difficulty. It is likely that the fixed sampling time of 500ms exceeds the true duration of evidence accumulation, a fact that could be probed by studying a reaction time version of this task, or reducing the sampling time.

## 2.5 Behavioral experiments reveal underlying decision strategy

All stimuli in this task included the same total number of dots across the screen. Visualizing the distribution of the stimuli used in this task along two axes (lower visual field strength vs upper visual field strength, Figure 2.3A) demonstrates that only a small proportion of all possible comparisons are presented to the animals on typical trials (black squares). Additionally, these trials lie across the diagonal, and as such, several discrimination axes would be equally accurate in producing the same choice behavior. For example, relying on the strength of only the lower visual field (Figure 2.3A, green decision axis), or only the upper visual field (not shown), would be equally effective in separating these stimuli and producing psychometric behavior as using the comparison strength between these two streams of information (purple decision axis).

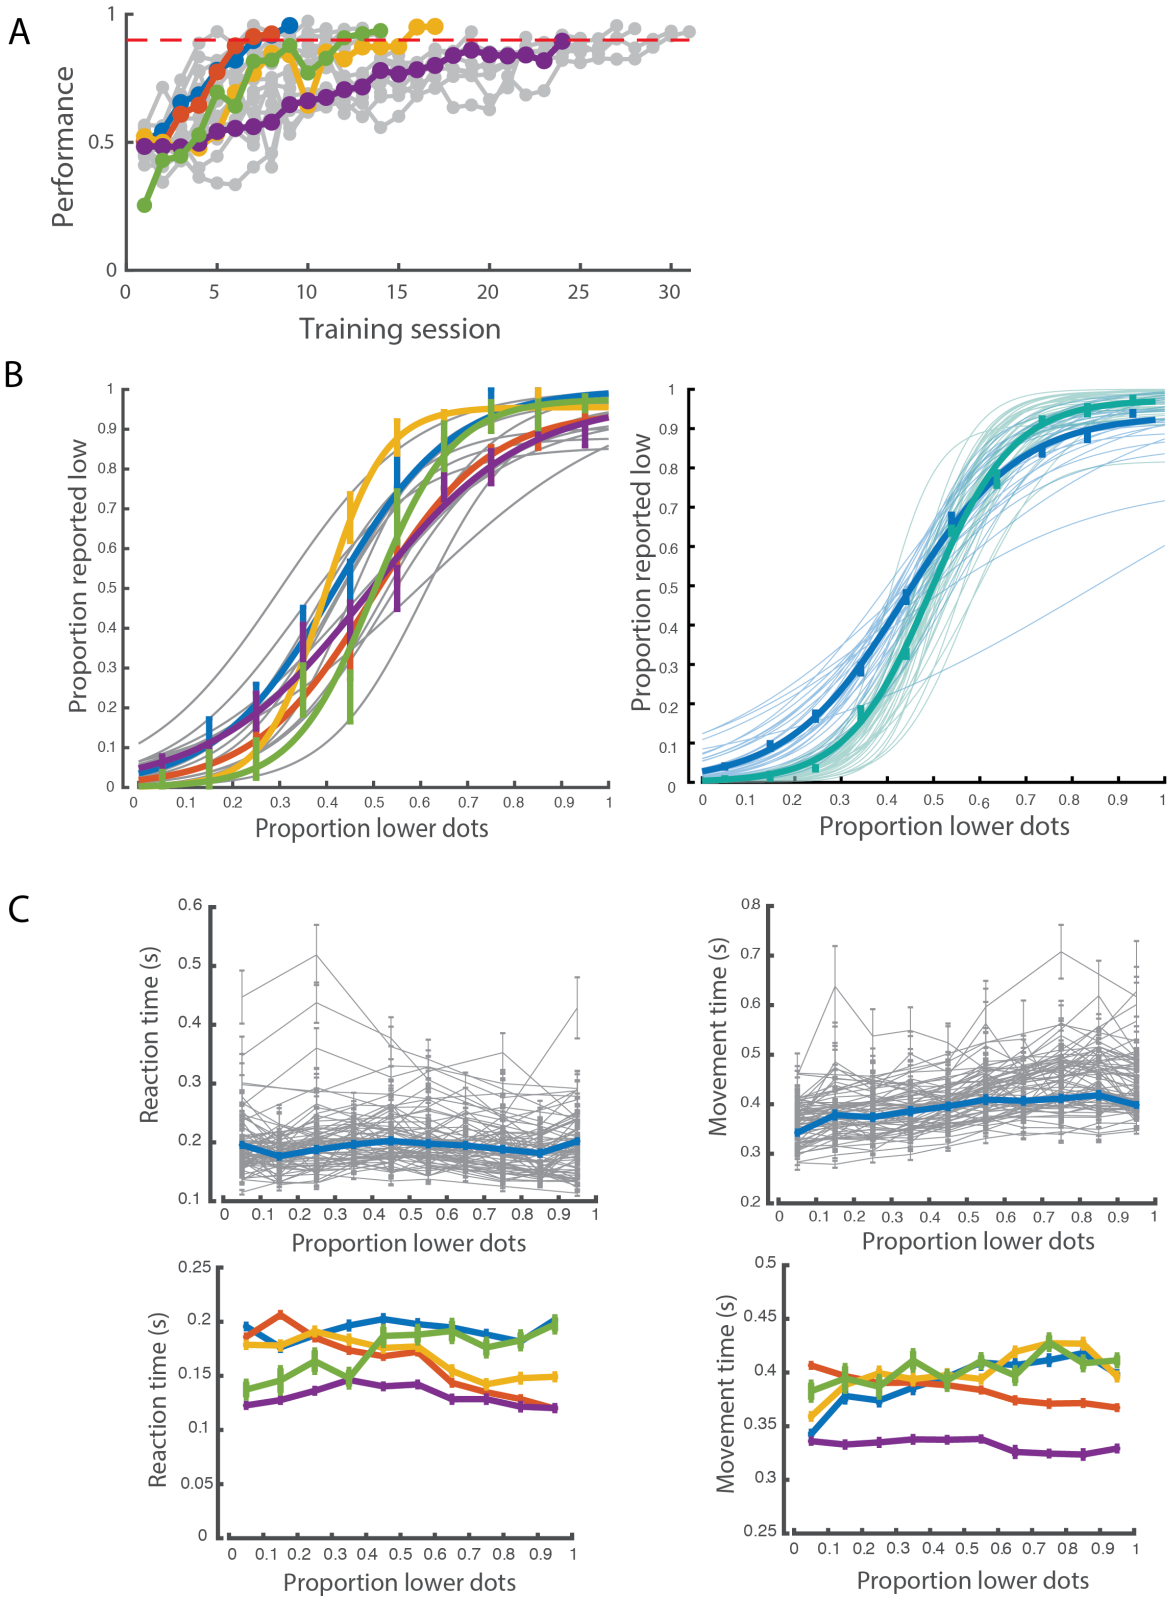
A
Performance
Training session
B
Proportion reported low
Proportion lower dots
Proportion reported low
Proportion lower dots
C
Reaction time (s)
Proportion lower dots
Movement time (s)
Proportion lower dots
Reaction time (s)
Proportion lower dots
Movement time (s)
Proportion lower dots

I wanted to understand which of these decision axes animals used to determine choice behavior. To probe this, I presented atypical stimuli to the animals on a subset of catch trials (5-10% of total trials) and measured the behavioral responses. Catch trials were randomly rewarded, to prevent learned associations during these sessions.

First, I presented stimuli wherein I maintained the same stimulus density in the lower hemifield compared to a typical trial for that stimulus condition, but the upper hemifield contained an equal density rather than the complementary (1-lower hemifield) proportion of dots (Figure 2.3A top panel). On these trials, there is zero comparison strength between the two subregions of the stimulus. However, along the axis of lower hemifield dot strength, these stimuli are identical to stimuli on normal trials, such that if one were to plot choice behavior against lower hemifield dot strength, an animal using only this stream of information would display the same psychometric behavior on these catch trials as on normal trials. Conversely, an animal using only the stimulus strength in the upper hemifield would display an inverted psychometric in comparison to normal trials when plotted against the lower hemifield dot strength axis. If the animal integrates both streams of information into a comparison value, the choice behavior on catch trials should be at chance level regardless of the stimulus strength in the lower hemifield. The choice behavior shown by two example animals is presented in Figure 2.3B. I found that unanimously, animals overweighted the lower hemifield stimulus strength at the expense of calculating a comparison metric between the two subregions of the stimulus. Indeed, animals were able to almost perfectly recapitulate their choice behavior using only the dot density on the lower half of the stimulus (Figure 2.3C). I noticed

FIGURE 2.2 *(preceding page)*: **Psychometric behavior on cloud of dots:** A. Learning curves: each line represents one animal. Colored lines represent animals later used for tetrode recordings. Color order corresponds to color order used in Figure 2.1C. Red dashed line represents 90% correct on easy trials. B. Single session psychometric curves. Left: for each animal in A, the psychometric curve on the last session shown in A, or if psychometric trials had not yet been introduced at time of 90% correct threshold crossing, the second session on which intermediate stimuli were presented. Right: Single trial and mean psychometric curves for an example animal prior to (blue - same color convention as in A, B left) and after (green) head position learning. Error bars indicate standard error of the mean.C. Chronometric curves for an example animal (top) and across animals (bottom). Neither reaction time (left) nor movement time (right) systematically varied with the stimulus difficulty on single sessions (top, grey) or across the mean (top, blue) for a single animal. Across animals (bottom), individuals had slight idiosyncratic differences in reaction time or movement speed depending on choice side, but not stimulus difficulty. Error bars indicate standard error of the mean.

a slight modulation at the upper end of the psychometric (i.e., when almost 2 times the typical total number of dots are presented on a given trial); this could be reflect either sensory noise due to higher overall illumination levels, or other proposed noise sources such as scalar variability (Scott et al 2015).

I confirmed these findings using a second type of catch trial where the lower hemifield signal was completely erased, but the upper hemifield ranged as it typically did between 5% to 95% of the total stimulus strength (Figure 2.3A,B middle panel). I wanted to ask whether, in the absence of stimuli in the lower hemifield, the upper hemifield strength would drive choice behavior. Instead, I observed that animals behaved identically across all of these trials, reporting the choice that corresponded to no dots in the lower hemifield, regardless of the stimulus strength in the upper hemifield. This is shown in Figure 2.3B (middle panel) where choice behavior is plotted against the "would-be" proportion in the lower visual field, which is present in control but not in catch trials. Indeed, animals report uniformly that there are more dots in the upper visual field on catch trials, independent of the magnitude of the difference between upper and lower hemifield strength on the catch trials, again supporting the hypothesis that animals solely base their choice on the signal in the lower hemifield.

Third, I used stimuli that had nonzero comparison value across the two stimulus subregions (Figure 2.3A, blue, orange, purple squares), but that were positioned off of the typical “diagonal” stimulus axis (black squares). Again, on these trials, I observed behavior consistent with thresholding of the lower hemifield strength, as choice behavior on these off-diagonal catch trials identically matched choice behavior on normal trials with the same lower hemifield density (Figure 2.3B,right panel).

From these experiments, I concluded that animals overwhelmingly and consistently preferred to use the information in the lower visual hemifield at the expense of relevant information that may be presented in the upper hemifield. I also found that the first type of catch trial (50% comparison trials, delivered at varying levels of stimulus densities) uncovered this strategy successfully across all animals, within single sessions. Therefore, moving forward, I used performance on this type of catch trial as a measure of decision strategy, and found that over 19 sessions across 6 animals, there was no single session in which animals did not recapitulate full psychometric behavior on these catch trials (Figure

2.3C). Because this strategy was absolute and uniform across all animals I tested, I hypothesized that this phenomenon is a species-level trait, rather than determined by individual preferences, and may therefore have some ethological relevance and/or neural circuit basis. The robustness of this preference across different stimulus distributions is explored in the next chapter.

## 2.6 Stimulus-independent variant of cloud of dots task

To enable comparison of visual processing of stimuli that are informative versus simply present in the environment, I designed a reduced version of the discrimination task where decisions were made stimulus-independent. In contrast to passive sensing paradigms commonly used in head fixed animals, subjects here are required to engage with the same task structure and motor actions as animals performing the discrimination task; the only difference is in the stimulus usage. I presented, on the majority of trials, a stimulus of randomly distributed dots across the full extent of the monitor. The number of dots on a given frame is the same as in the discrimination task. On a subset of trials, I presented one of the two easy stimulus conditions from the discrimination task (5% dots in lower half or 95% dots in lower half). Following successful fixation through the stimulus period, animals were rewarded for a choice report to either choice port.

One important difference between this non-visually guided behavior and behavior on the discrimination task is that following learning that reward delivery was stimulus-independent, and further, that both ports were equally viable at all times, animals converged over time onto a single preferred reward port. Therefore, while on the discrimination task both choice directions were sampled equally, on this reduced version of the task, animals overwhelmingly made choices in a single direction. For some animals this side preference was complete, while for others, sampling of the non-preferred port happened on between 1-5% of trials (Figure 2.4B). This therefore limited the analyses that could be applied to neural data recorded on this version of the task. However, I chose not to implement any history- or stimulus-dependent rules to force sampling at both reward ports, as this would further deviate from the

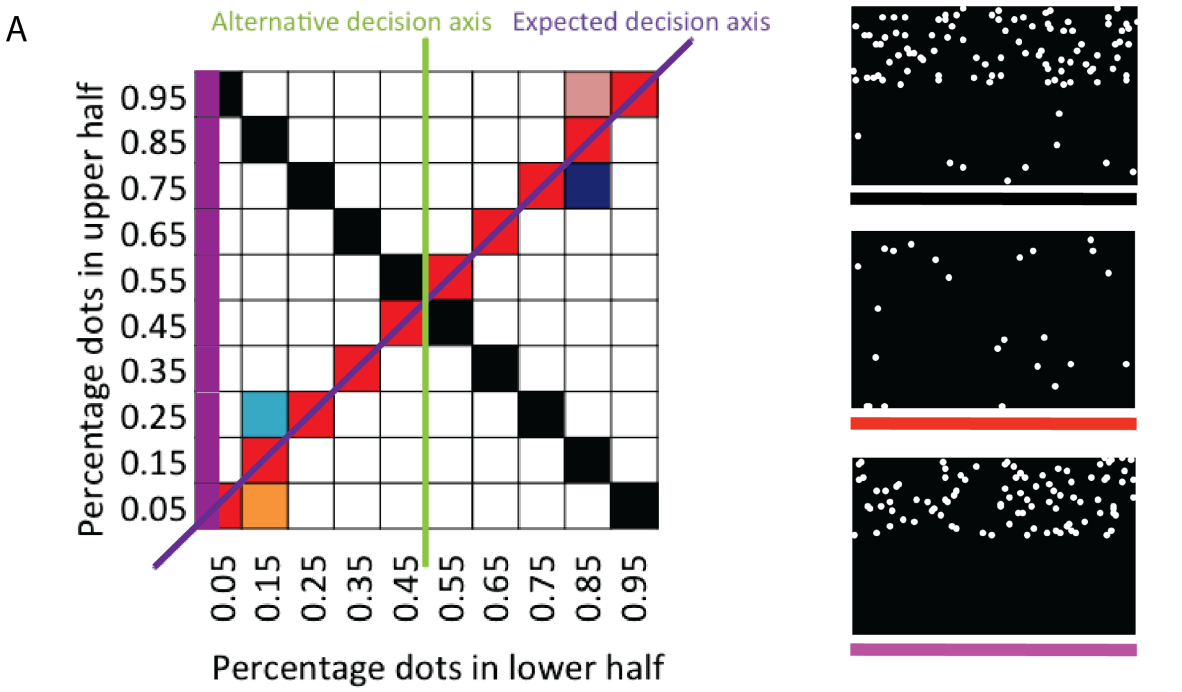
A
Alternative decision axis
Expected decision axis
Percentage dots in upper half
0.95
0.85
0.75
0.65
0.55
0.45
0.35
0.25
0.15
0.05
0.05 0.15 0.25 0.35 0.45 0.55 0.65 0.75 0.85 0.95
Percentage dots in lower half

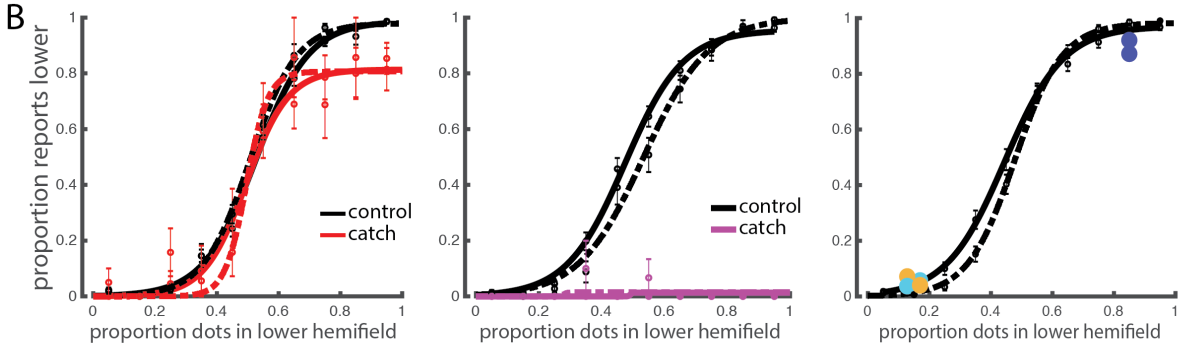
B
proportion reports lower
control
catch
proportion dots in lower hemifield
control
catch
proportion dots in lower hemifield
proportion dots in lower hemifield

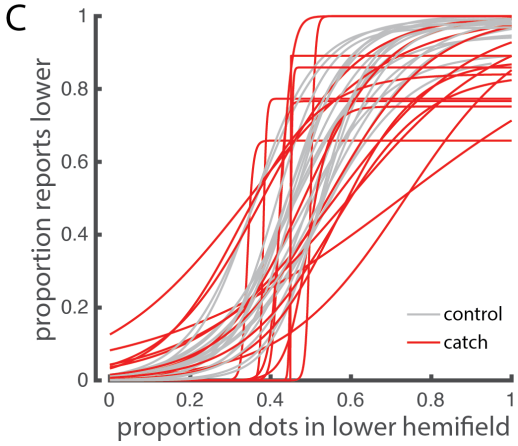
C
proportion reports lower
control
catch
proportion dots in lower hemifield

structure of the discrimination task, and introduce additional task parameters that may be reflected in V1 processing.

The other major way in which this task deviated from the discrimination task was in the absence of a punishment outcome. However, despite reward certainty on all completed trials, animals still failed to collect reward on a small percentage of trials (Figure 2.4C). The proportion of missed trials over a session varied between 0-20%.

Moving forward, I use this task structure to interrogate which features of V1 activity patterns are dependent on stimulus usage for task execution. I will explore the nature of these responses in Chapter 4.

## 2.7 Methods

I used the Bpod system (Sanworks, NY) to implement the behavioral state machine. The task structure was as follows: animal entry into the center port triggered the beginning of a prestimulus delay. For each trial, I drew a variable pre-stimulus delay from an exponential function with a mean of 0.3s. Following this delay, the stimulus presentation was triggered through Psychtoolbox (Brainard, 1997, more below).

FIGURE 2.3 *(preceding page)*: **Catch trials reveal behavioral strategy:** Responses to atypical stimulus conditions reveal underlying behavioral strategy. A. Left: "Map" of stimulus conditions on control trials (black squares), and different types of catch trials (colored squares), plotted by the relative stimulus density delivered in the upper and lower hemifield. "Proportion of dots" gives the strength in a given hemifield relative to the total number of dots delivered over the full screen on control trials. The absolute strength in the lower subregion of the stimulus for a given trial can be read off from its position along the x axis, and the absolute strength in the upper subregion from its position along the y axis. The comparison strength for each of these trial type is equal to the difference between its x-coordinate and y-coordinate on this plot. All measurements are in units of percentage of total dots on original trials. Right: Example stimulus frames for a control trial (top, black label), 50% comparison trial (middle, red label), and hidden lower hemifield trial (bottom, purple label). Each of these frames corresponds to the same "lower hemifield strength", and will be plotted at the same position on the x-axis in following psychometric plots; in the red trials, the upper hemifield has been manipulated to match the lower hemifield, while in the purple trial, the lower hemifield has been erased, and only the complementary upper hemifield remains. B) Psychometric behavior, pooled across sessions for two example animals. Black lines indicate control trial behavior, while colored lines and points indicate behavior on catch trials. Error bars indicate standard error of the mean. C) Single session psychometric behavior on the 50% comparison catch trials (N=6, n= 19). Grey lines indicate control trial behavior, red lines indicate catch trial behavior.

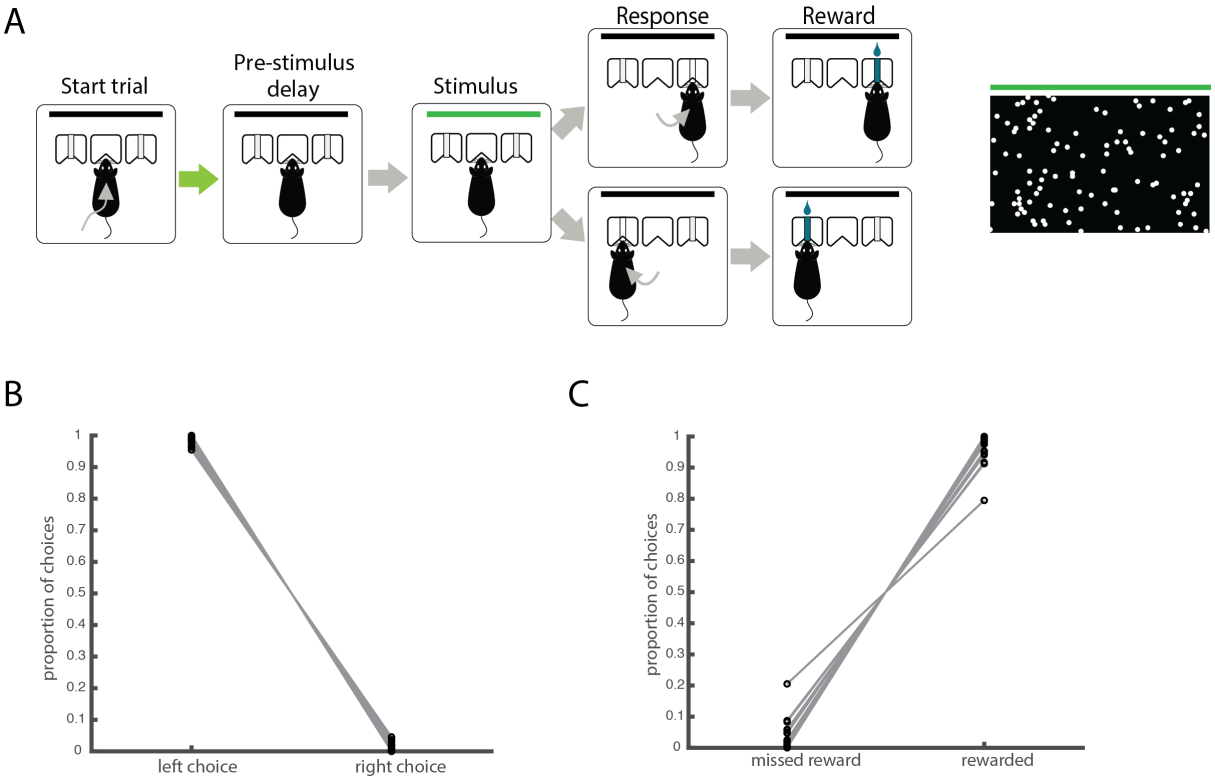


FIGURE 2.4: **Passive viewing task design and behavior:** A. Task design. Green arrow reflects the same head positioning procedure as in Figure 2.1.

I measured the delay for Matlab and Psychtoolbox to trigger stimulus onset using a luminance reporter module (more below), and found that onset was delayed by 107ms on average. To account for this, I added a 200ms fixed post-stimulus delay between the stimulus off trigger and the decision tone. Any withdrawal from the center nosepoke at any point between the pre-stimulus delay initiation and the decision tone delivery led to a missed trial and a 2s time out. Upon learning of the head position protocol, a missed trial could also be triggered by a head movement while in the center port during this peristimulus period. After the decision tone, the animal was allowed to withdraw from the center port, and make a choice by poking into one of the two side ports within 3s. A 20 μL reward was delivered following a 50ms nosepoke into the correct port. A correct choice report that did not fulfill this duration requirement did not trigger reward, but no punishment was delivered either. No intertrial interval was specified following correct (either rewarded or missed reward) trials. If the animal made an incorrect choice, a punishment tone (white noise stimulus) was delivered, as well as a 5-8s time out. The length of the time out was typically 5s long, but for some animals this was increased up to 8s to discourage incorrect choices (or lapses) on easy trials.

I used the Psychtoolbox set of Matlab functions to generate and deliver both our visual stimuli and our auditory decision and punishment tones. For each stimulus, I generated 30 frames, to be delivered at 60Hz, wherein stimuli were randomly distributed across each frame following the stimulus condition on that given trial. I defined two subregions of equal size, separated by a thin boundary region where no dots were ever present. To determine absolute dot numbers in each subregion on each frame, I drew values from a Poisson distribution centered on the lesser value of the stimulus condition. For example, a trial that consisted 25% upper and 75% lower stimuli delivered on any given frame $n$ points drawn from the Poisson distribution centered on 25% of total dots; the number of dots in the lower half of the screen was then set to *total-n* points. Therefore, every frame had the same total number of individual dots. Though these draws meant each individual frame of a stimulus had a slightly different proportion of dots across the two subregions, I was able to recover the true mean of the stimulus condition by averaging across frames. Across all frames, I set the total number of dots such that the density across the screen was 1% of all possible locations. At each dot location, I presented a round white dot that subtended about 3° in

visual space. While this is smaller than the mean receptive field size in rodent V1, I expected that the spatiotemporal structure of this stimulus would necessitate integration of individual stimuli over both space and time. I generated all stimuli for all trials at session startup, then used Psychtoolbox, triggered through Matlab, to display the pre-loaded stimulus at the time of stimulus onset.

I generated catch trial stimuli similarly, but adjusted the absolute number of dots in each subregion according to the type of catch trial that was being presented. For example, for 50% comparison catch trials, instead of delivering the complementary number of dots in the opposing subregion, I delivered the same number of dots. On catch trials, the number of total dots, and therefore the total luminance, across the full screen was not consistent with either typical trials or other catch trials. Catch trials were rewarded randomly regardless of choice side.

Stimuli for the stimulus-independent task were generated similarly, except that individual dots were truly randomly distributed across the monitor. For the majority of trials, there was no boundary region. However, for the subset of trials (usually about 20% of trials) on which I presented the easy stimuli from the cloud of dots task, the boundary region was defined identically to in the stimuli from the cloud of dots task. In this task, the stimulus period was increased to 700ms, so more frames were drawn (42 frames) for each trial.

I adapted a luminance detector module (Frame2TTL, Sanworks) to report luminance changes during each trial and to signify the onset of the stimulus delivery. I placed this module above a reporter pixel which flickered on/off with each frame update. I also used this module to check the consistency of the luminance of the monitor over time.

Finally, I implemented the closed-loop head position condition using Bonsai, a reactive programming software (Lopes et al., 2015). Bonsai was given video input from a webcam mounted above the animal at a 70° angle. This video input was binarized and regions of interest (ROIs) were defined on a per-animal basis from this field of view. These ROIs were centered on the position of each ear, such that the ear would entirely fall within the ROI when properly aligned. Then, I used built-in Bonsai functions to implement contour mapping of the image within each ROI, and determine the size of the object. I

filtered for object sizes of a certain range, and then calculated the x and y position of the centroids of these objects. I compared the distance between the two pairs of x and y values to a threshold to generate a binary signal representing the animal's alignment. Further, to ensure I only tested for this condition when the animal was in the port, and therefore prevent any spurious detections caused by background (e.g. behavior rig floor) objects, I sent an input trigger to this algorithm from Bpod only when the animal was poked into the center port during a valid trial. The algorithm then performed a moment-to-moment logical AND computation on the comparison between the x values, the comparison between the y values, and the input trigger to output a binary trigger back to Bpod. The continuation of the Bpod states depended on the continuous on-state of this trigger. To ward against fast software- or camera-generated errors from producing too many false negatives, I implemented a short 50ms grace period that was entered whenever the trigger turned from on to off. If during this grace period the trigger returned to the on state, the trial was allowed to continue; otherwise, it was aborted.

Psychometric curves were generated using the Matlab function `fit` to fit a logistic function to single session or across-session pooled choice behavior:

$$y = \frac{l}{1 + exp(-beta * (x - alpha))}$$

where the fit parameter alpha approximates the bias term of the psychometric function, and the beta term approximates the slope of the curve. These parameters will be used in later chapters to compare psychometric curves across sessions and conditions.

## 2.8 Discussion

Here, I described a novel visual discrimination task for freely moving rodents, where I additionally harnessed video processing technologies to implement a virtual head fixation protocol that is both non-invasive and compatible with future experimental techniques. I showed here that animals were able to

learn this head position rule without any direct reinforcement signal (e.g. punishment tone upon misalignment), which further increases its compatibility within a given task as it requires minimal changes to the trial structure. This could be useful to implement in tasks that require either spatial discrimination (e.g. sound localization tasks) or whose stimulus depends on spatial location more finely (e.g. sound frequency discrimination). While animals are proficient at learning and sustaining this type of fixation, physical head fixation methods, either sustained or transient (e.g. Scott, Brody, and Tank, 2013), will be necessary to allow the adaptation of this task for applications requiring absolute stability such as imaging techniques. Recently developed deep-learning pose estimation algorithms (Mathis et al., 2018) that process images at high resolution can be used to quantify the angular change in head position over learning and within sessions; however, it remains difficult to implement online head position tracking using these algorithms.

I showed that rats systematically preferred to make visual decisions using information available in the lower, rather than upper, visual field. What remains unclear from these data is whether rats are bound to always use the lower visual field for visual decisions, or if the observations here are due to rats learning from the stimuli sampled over training that in this case, a single-stream thresholding strategy is an equally (or perhaps, if independent noise of each information stream is taken into account, more) successful method to solve this task. However, if a thresholding strategy were suboptimal for the task at hand, would rats be able to implement an alternative strategy? What might be the neural basis for the observed bias, and is it absolute or flexible? I will explore this question further in Chapter 3.

Many questions remain regarding the decision process in this task. I have implemented a fixed time decision task, but analysis of response times and movement speeds between ports suggest animals are not operating at the limit of their decision process. The lack of scaling of response latencies with stimulus difficulty suggests that animals have, at 500ms of stimulus sampling, sufficient information to make their decision, and require no further time for deliberation in the absence of further sensory input. Conversely, any effects may be obscured by our use of a fixed stimulus duration and decision tone to signal to the animals when a decision is available; it is possible that more time could be taken if animals were not

trained to know that additional time would not help their choice, as no additional stimulus information is available, and may in fact reduce their rate of reward. A reaction time task would be an easy extension to probe the temporal limits of visual spatial discrimination in rats.

This task was strongly inspired and adapted from an analogous auditory task. The similar task structures and analogous stimulus dimensions will allow us to ask cross-modality questions in the future, such as: how do representations across primary sensory cortices compare with respect to sensory- and non-sensory encoding? Do different sensory modalities share analogous output circuitry when behavioral demands and stimulus complexity are comparable? Put another way: how much of goal-directed sensory processing is based in shared principles and pathways, and how much is specific to individual sensory modalities?

# Chapter 3

# Behavioral strategy is dynamic over training

## 3.1 Introduction

In the previous chapter, I described our finding that rats display a consistent and seemingly intrinsic bias in our task to use the information in the lower stimulus subregion rather than either the upper subregion or a comparison between the two subregions. I wondered whether this phenomenon was due to the specific task design, and to what degree it would affect a task in which the two subregions of the stimulus were not always complementary and therefore informationally redundant with one another. To explore this question and further characterize the basis and limits of this intrinsic bias, I trained and tested rats on two variants of this task where I delivered a wider diversity of stimulus draws, and observed how the rats' strategies evolved with time and stimulus distribution.

Many psychophysical tasks require the subject to judge the greater of two quantities (Uchida and Mainen 2003, Znamenskiy and Zador 2013, Brunton et al 2013, Akrami et al 2018). It has been argued that animal strategy can in many of these cases be constrained by varying the relevant stimulus parameter widely across trials, or implementing "2-dimensional" stimulus designs (Romo et al 1999, Brunton et al 2013, Akrami et al 2018). In these cases, solely the information in one part of the stimulus (e.g. number of left auditory clicks, or stimulus volume) should not be sufficient to produce accurate choice behavior,

whereas when the two stimulus streams bear a consistent relationship to one another (as in a mixture of two stimuli, e.g. odors), it may be possible for subjects to find a different solution, as they do in the task I described in the previous chapter. Here, I show that even in cases where two stimulus streams are independent of one another, alternative strategies can be used to solve the task along a different axis from what is intended, and this may be undetectable from simple psychometric curves.

## 3.2 Rats overutilize information in the lower visual field at the expense of optimal reward collection

The first variant of the cloud of dots task was similar to the original task design (Figure 3.1A, left) but included additional "off-diagonal" trials (Figure 3.1A, center, blue and green squares) at a proportion of 70-90% of all trials. The stimuli on these additional trials were drawn such that the number of dots across the two stimulus subregions summed to either 60% (green) or 140% (blue) of the total number of dots presented on the original trials (black) (Methods). The comparison strength for each of these conditions is equal to the difference between the x-coordinate and y-coordinate of a given trial type (square). I hypothesized that increasing the exposure to trials that did not always sum to the same total stimulus strength would deter animals from finding redundancy in the two information streams, and would compel them to gather and compare information across both information streams (upper and lower).

I found that while animals initially were able to treat all three types of trials (original, more dots, and fewer dots) similarly relative to the strength of the comparison between the upper and lower subregions (Figure 3.1B, top), over time, they converged again on a thresholding strategy, whereby the lower stimulus strength was sufficient to explain all choice variability across all three trial types (Figure 3.1B, bottom). This was the case despite the fact that the two decision axes no longer gave the same decision accuracy rate, as they did in the original task.

I further extended this task design to sample uniformly across the full stimulus distribution (i.e.

independent draws of dots in the upper hemifield vs lower hemifield, Figure 3.1A, right). Again in this task design, the two candidate decision strategies do not give the same rate of reward. Specifically, for an ideal observer presented with a uniform distribution across this space, a thresholding strategy will result in only 75% accuracy while a comparison strategy will result in 100% accuracy. However, subjects again uniformly converged on a thresholding strategy over training, at the expense of this 25% drop in reward rate (Figure 3.1C, bottom). I saw this trend uniformly across all animals trained on this task (N = 5).

## 3.3 Rats adapt their task strategy based on statistics of stimulus distribution over trials

In light of the finding that preference for using solely lower hemifield information was robust even when such a strategy was suboptimal, I wanted to understand the conditions under which animals would be able to stably carry out a comparative task strategy. From the behavior profiles during early learning, it was clear that information from the upper hemifield of the stimulus could be integrated into a visual decision, but animals simply chose not to continue doing so after extensive exposure to the task. I noticed that in the uniform distribution task, certain subregions of the stimulus distribution were more informative than others about which strategy was optimal. In the schematic above the arrow in Figure 3.2A, stimuli in the grey regions are equally correct for an ideal observer using a comparative strategy and an ideal observer using a lower hemifield thresholding strategy. However, the stimuli in the black regions of this schematic will give opposing results under these two strategies. Therefore, I wondered if behavioral strategy could be shaped by enriching the proportion of stimuli residing in the black regions over trials (here termed adversarial trials).

I found that indeed, animal strategy was susceptible to the statistical structure of the stimulus distribution over trials. Comparing choice profiles over the full stimulus space before and after adversarial training, I find that choice proportions became dependent on both the lower hemifield strength and upper hemifield strength following adversarial training (Figure 3.2A). Psychometric behavior on trials with more and fewer total dots also became more similar to each other when aligned to comparison strength

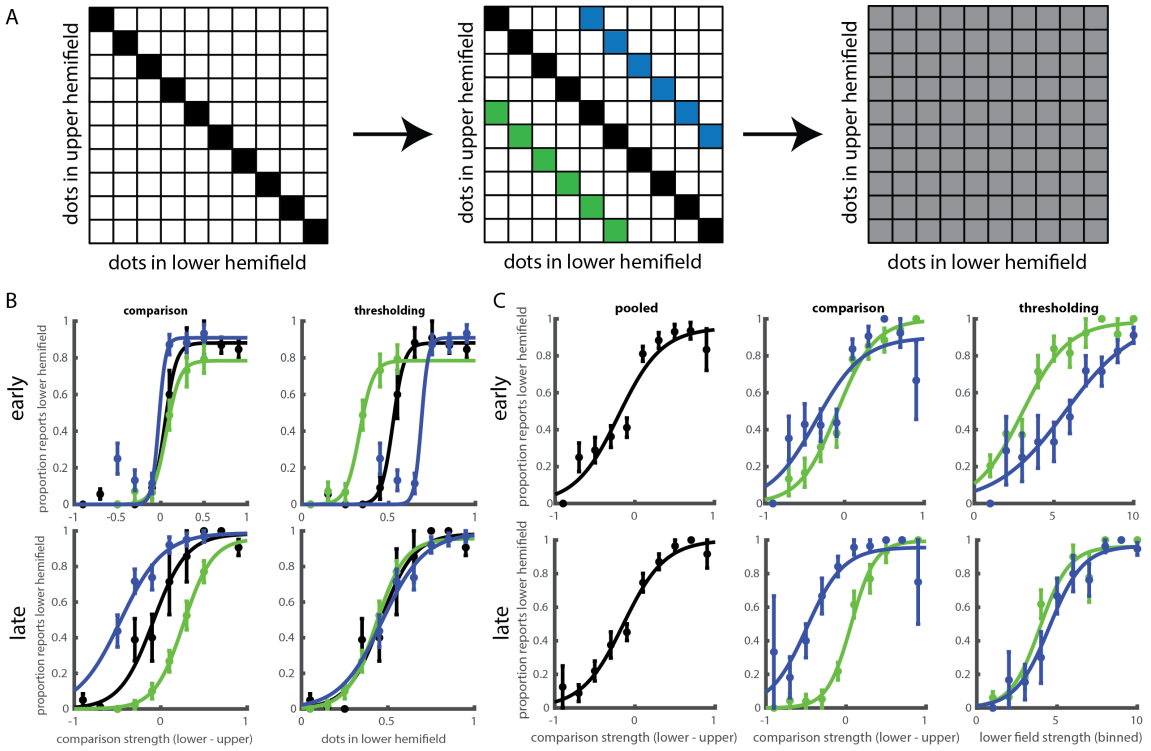


FIGURE 3.1: **Rats converge on lower visual field thresholding strategy over training:** A. Changes in stimulus distribution across the variants of the task. Left: original task, middle: off-diagonal stimuli added, right: uniform draw across full stimulus distribution. All axes are in units of percentage of total dots on original trials. B. Strategy evolves with training time on task variant with original and off-diagonal stimuli. Black lines describe behavior on original trials (complementary upper and lower subregions), while blue lines describe behavior on trials with stimuli summing to 140% of original dots (more dots trials), and green lines describe behavior on fewer dots (60% of original dots) trials. Psychometrics are plotted either along the axis of comparison strength (lower - upper) to test for use of comparison strategy, or lower hemifield dot strength to test for use of thresholding strategy. When an animal is using a given strategy, I expect the three psychometric curves to align along that axis, meaning that stimulus axis explains all of the variance in the choice behavior, and the total number of dots across the screen does not provide extra information. Top and bottom rows correspond to single sessions at different points in the training process. Early on in training, the psychometric curves largely align when plotted along the comparison strength axis, while later in training, after 1-4 weeks, psychometrics curves are aligned relative to the lower hemifield dot strength. C. The same as in B, but for the uniform draw variant of the task. Here, black trials represent all trials, pooled, blue trials represent all trials with more dots than the mean number (original number of dots), and green trials represent all trials with fewer dots than the mean number. Similarly to B, task strategy moves from comparison to lower hemifield thresholding over training.

(Figure 3.2B). Meanwhile, lower visual field strength failed to explain choice behavior. Most importantly, performance on adversarial trials, when plotted by the comparison strength, matched performance on non-adversarial trials following training (Figure 3.2B, right). Moving forward, I used the normalized performance on adversarial trials as a measure of how well an animal carried out a comparison strategy:

$$performance = \frac{P_{correct}adversarial}{P_{correct}non-adversarial}$$

I used this measure to understand the timecourse of strategy change during adversarial training. I first increased the proportion of adversarial trials twofold, from their baseline 25% proportion of all stimulus to 50%. Performance on adversarial trials, normalized by the accuracy on all other trials, increased only modestly following this change (Figure 3.2C). However, when I increased the proportion of adversarial trials to 75% (a threefold increase from baseline), every animal I trained was able to perform as accurately on adversarial trials as on non-adversarial trials.

Finally, using this metric, I tracked the behavioral strategy for 3 weeks after returning the stimulus distribution to a uniform draw following a period of adversarial training. All animals tested were able to maintain high accuracy on adversarial trials relative to non-adversarial trials for over 3 weeks (N=3). Notably, two animals started off with higher accuracy on adversarial trials relative to baseline (normalized adversarial performance > 1), suggesting possible overfitting to the adversarial stimulus distribution following extended training. This overfitting effect dropped off quickly to the unity level within 3-4 sessions, but performance on adversarial trials did not continue to drop substantially after this, indicating animals stably carried out the comparison strategy for weeks even after the stimulus statistics that forced that strategy were relieved.

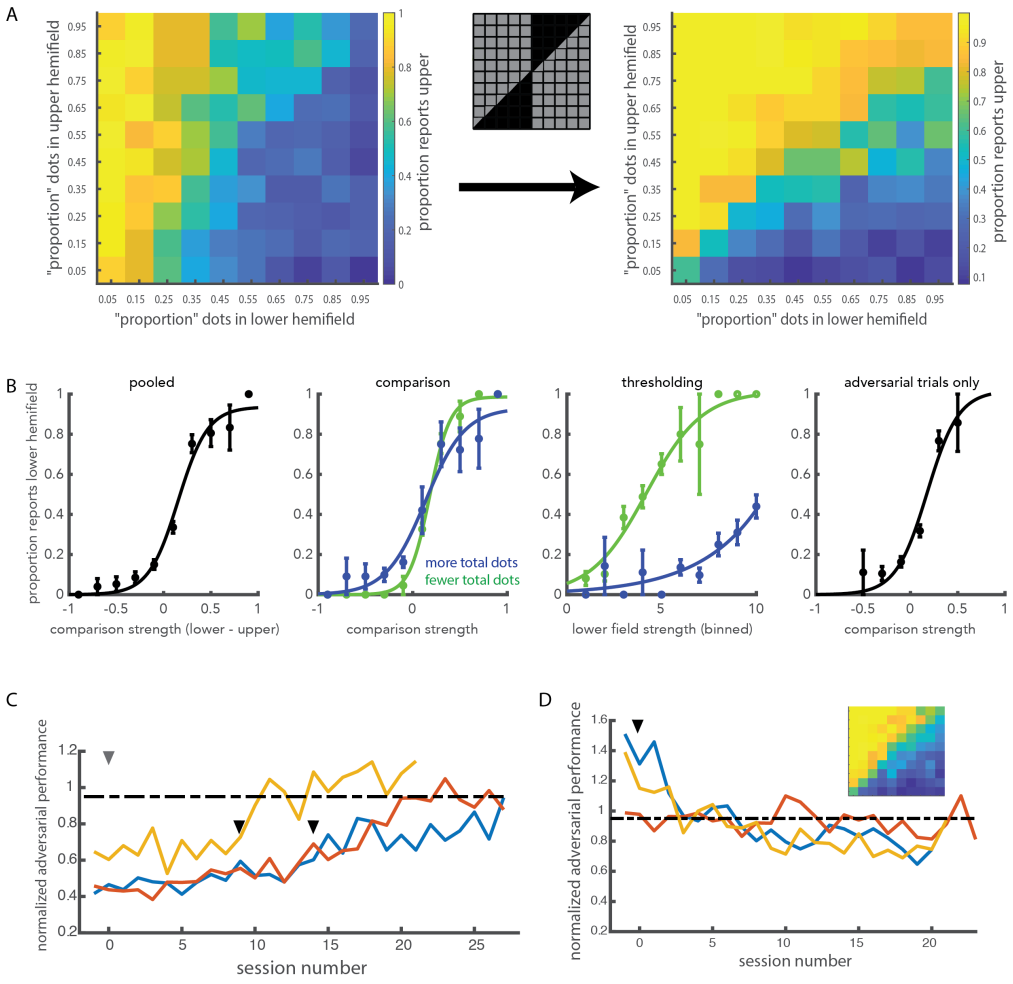


FIGURE 3.2: **Behavioral strategy shifts with the statistical structure of the sensory environment:** A. Heat plots showing behavioral profiles across the 2D stimulus distribution before and after adversarial training. In the heat plots, each square corresponds to a binned trial condition whose mean lower visual field strength and mean upper visual field strength are denoted on the x and y axis, respectively. The color of each square corresponds to the choice behavior of the animal in response to the stimuli in that bin. Warmer colors correspond to the animal reporting that there are more dots in the upper visual field; cooler colors mean more reports of more dots in the lower visual field. Each heat plot is built from 4 consecutive behavioral sessions. The heat plot on the left shows the behavior across 4 sessions prior to adversarial training. Note the non-uniform choice profiles across trials with the same comparison strength (diagonals from lower left to upper right). The heat plot on the right shows behavior across 4 sessions from the same animal following adversarial training. B. Psychometrics on a single session from the same animal in A, demonstrating that behavior is consistent with implementing a comparison rule following adversarial training. The left three plots follow the same conventions as in Figure 3.1C, while the rightmost figure shows performance only on adversarial trials, plotted against comparison strength. In this example session, adversarial trial behavior matches or slightly exceeds baseline pooled behavior (left panel). C. Strategy change over adversarial training can be seen from animals' (n=3) improved performance on adversarial trials relative to baseline. Grey arrow denotes session when adversarial trial proportion was increased to 50% (day 0). Black arrows denote session when adversarial trial proportion was increased to 75% (day 9 for yellow curve, day 14 for blue and orange). Black dashed line denotes 95% accuracy relative to baseline trials. D. Strategy trajectory after stimulus distribution was returned to uniform draw post-adversarial training. Black dashed line denotes 95% accuracy relative to baseline trials.

## 3.4 Asymmetric V1 corticostriatal projections correlate with observed behavioral bias

The circuit basis allowing V1 activity to drive downstream choice in this task is not known. Based on work in an analogous auditory task, I was curious about the possible role of the corticostriatal projection as an output pathway from primary visual cortex. Therefore, I sought to map this anatomical projection using viral tracing tools.

A direct projection from V1 to striatum has previously been reported and characterized in mice (Khibnik, Tritsch, and Sabatini, 2014). However, one crucial feature of the auditory corticostriatal projection was the preserved topographic organization of these projections, such that localized regions of the posterior tip of the striatum preferentially received either low-frequency preferring projections or high-frequency preferring projections (Xiong, Znamenskiy, and Zador, 2015). I was curious to see if such an organization was also present in the visual cortex projection to striatum. I therefore applied dual fluorophore tracing to visualize these projections, targeting two locations along the anteroposterior axis of V1, which correlates with the elevational retinotopic axis.

What I observed, however, looked very different from the auditory projection pattern. While auditory corticostriatal neurons densely innervate the posterior striatum, V1 projection neurons terminated only very sparsely in the posterior tip of the striatum (Figure 3.3C right panel). Instead, V1 projection neurons tended to terminate more strongly along the dorsomedial extent of the mid-posterior striatum (Figure 3.3C left, middle panels). More importantly, I noticed that these terminations were not equal in strength from the two injection sites. Across animals, I found that the posterior injection site sent fewer axons to striatum; in no striatal section did the terminations from the posterior injection site outweigh the terminations from the anterior injection site. This was in stark contrast to terminations in superior colliculus, where termination strength was roughly equal across the two injection sites (Figure 3.3B). Further, while the terminations in superior colliculus showed a clear axis of separation, therefore maintaining the retinotopic separation at least along one dimension, the terminations in striatum were much less clearly separated.

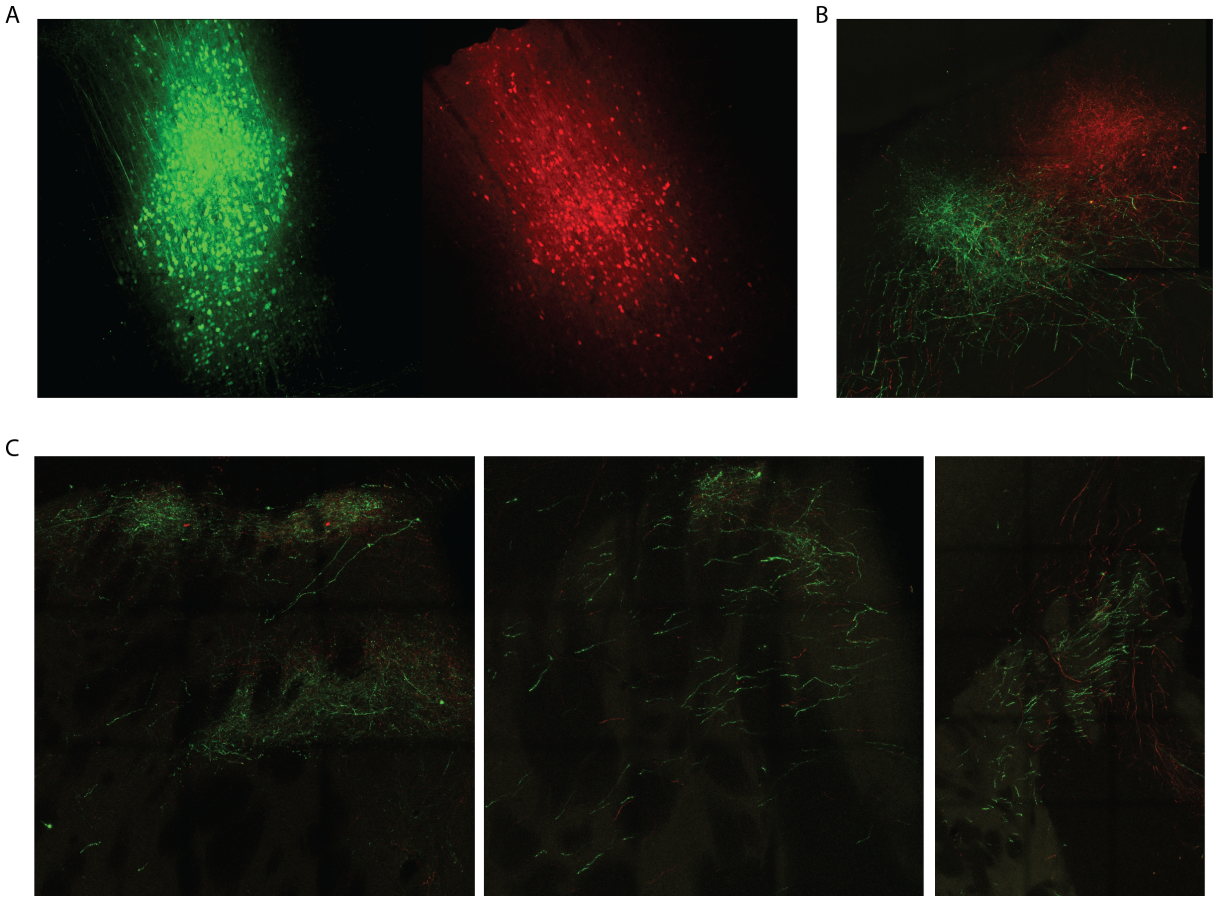


FIGURE 3.3: **Asymmetry in projections from V1 retinotopic subregions to striatum, but not superior colliculus:** A. Injection sites in anterior (bregma - 5.5mm AP, +4.5mm ML, left) and posterior (bregma -6.5mm AP, + 4.6mm ML, right) primary visual cortex. B. Terminations from cell bodies in A in ipsilateral superior colliculus. C. Terminations in ipsilateral striatum, from anterior to posterior.

Given the importance of the corticostriatal projection in the auditory task, and the correlation between the weak upper-visual field preferring projections to striatum and the robust bias against upper-visual stream information in several instantiations of the cloud of dots task, this projection may be a good candidate for further study. It would be particularly interesting to understand whether this anatomical structure plays a role in determining or constraining the decision strategy, manifesting as a behavioral bias toward the lower visual field.

## 3.5 Methods

All tasks in this chapter shared the same task structure as the original task described in Chapter 2. To add "off-diagonal" trials to the original version of the task, I drew the number of dots for the lower subregion as a percentage of the total dots on the original trials (e.g. 5%, 15%, 25%, 35%, 45%, 55%, 65%, 75%, 85%, 95%) and then placed in the upper subregion the number of dots that would bring the total across the full screen to the total for that type of trial (60% for "fewer dots" trials, and 140% for "more dots" trials). To uniformly cover the stimulus space, I randomly drew independent values between 0 and 1 for the two subregions, and applied these as scaling factors to the total number of dots in the original trials, to generate the number of dots to be presented in each subregion.

To increase the proportion of adversarial trials, I checked the position of all stimulus draws, and determine the number of trials needed to bring the adversarial trial proportion to the specified level (50-75%). On those trials, I maintained the trial category (more upper or more lower) and re-drew values for the upper and lower field that would move the trial into the adversarial region. For example, if a trial was classified as more lower, I generated a stimulus where the proportion of lower dots was less than 0.5, but the proportion of upper dots was even less, between 0 and the currently selected proportion of lower dots.

To label projections from V1, I injected anterograde viruses (AAV1.CAG.tdTomato, 15nL/injection site, and AAV1.CAG.GFP, 20nL/injection site) into either frontal (bregma -5.5 – -6mm AP, +4.5mm ML) or slightly more posterior V1 (bregma -6.5 – -7mm AP, +4.6mm ML). I injected at 800 μm below the surface of the brain, to avoid any labeling extending below cortex. After waiting at least 14 days for viral expression, animals were sacrificed and brains were collected. Vibratome sections were mounted and imaged at 20x.

## 3.6 Discussion

From this set of experiments, I learned that animals' preference for using the lower visual field to drive discrimination decisions persisted even when such a strategy was suboptimal, and was not simply the result of redundant stimulus design. I also learned that animals tended to sample information streams more broadly earlier in training, such that they were able to derive and discriminate based on a comparison value, but eschewed this strategy over time and exposure to the task. The reason for this is unknown, but there are several possibilities. The first is that integrating across two information streams may lead to an additive effect on behavior of the sensory noise in the estimates of each of those information streams, and the decrease in reward rate generated by selecting the simpler strategy is offset by the decrease in sensory noise. A second possibility is that there is some unknown "effort" variable that enters into the equation, whereby animals are willing to sacrifice absolute levels of reward for simpler strategies.

That animals converged on this alternative strategy in the uniform distribution task was surprising for two reasons. First, while the original task delivered stimuli that bore a consistent relationship between the two subregions, the two streams of information in the uniform distribution task are by definition independent of one another. Therefore, it is unlikely that a shift to a thresholding strategy in this task is due to an awareness about stimulus structure on a per-trial level. Second, together with the fact that they sampled the more optimal strategy early in learning, this suggests that animals may integrate stimulus structure across trials to find a less, not more, optimal strategy. While other studies have rightly noted that in 2D tasks like this, which vary independently across two dimensions, animals cannot optimally solve the task along only one of those dimensions, it appears here that they can still try, and do quite well. In fact, it is difficult to see from only the psychometric curve that they are implementing a suboptimal strategy. One possible reason why animal carry out this suboptimal strategy in this particular task is that stimuli do not vary widely enough across each of these dimensions. There are intrinsic limits to this, given how the stimuli are delivered on the screen. For example, if there are zero stimuli on the lower half of the screen, chances are very good that there will be more dots than that in the upper subregion. At the very least, both regions will be zero, so a "more-upper" report will necessarily have a greater than 50%

probability of reward. As animals are exposed to many trials over learning, the boundaries of the stimulus distribution may become more clear, and thereby influence the basis of single trial discriminations.

Along similar lines, I also observed that animals can be pushed back into stably carrying out a comparative strategy by changing the statistics of the sensory stimuli they are asked to judge. When I changed nothing about the stimulus structure on individual trials, but only the relative frequency with which trials distinguishing the two candidate strategies were presented, animals were consistently able to learn to carry out the strategy that earned them the most reward. What remains unknown is which feature of the task environment animals used to drive this learning. One possibility is that animals track the overall decrease in reward rate (up to a 12.5% difference in reward rate over the course of adversarial training) and at a certain threshold categorically change their information use. Another possibility is that animals build an expectation of the stimulus distribution over trials, and adjust their strategy as that expectation changes. It will be interesting to understand in the future which of these mechanisms comes into play to drive behavioral flexibility, even in overtrained behaviors such as those presented here.

Other open questions include why execution of the comparison strategy remains stable for >3 weeks following adversarial training, but changes more quickly during initial learning of the task. How does readout of the two information streams adapt in neural circuitry when categorical shifts happen, and a given stimulus subregion is ignored or brought back online? Finally, especially during adversarial training and sometimes during early training, behavior appears to rely on combination of these strategies. Given that each trial can only result in a single binary choice, are these strategies combined within or across trials?

The experiments in this chapter point primarily to the influence of the stimulus distribution over trials on single trial discriminations. An important related point is how difficult it is to see the true decision axis of an animal when behavior across two dimensions is collapsed down into a single psychometric curve. This is more likely when two possible decision axes are more similar to one another, but can affect any case in which they are non-orthogonal. However, splitting the data along different dimensions is an effective method to isolate the stimulus feature that makes the greatest contribution to choice behavior.

# Chapter 4

# Neural representations in behaving animals

## 4.1 Introduction

The classical view of early sensory cortices is that their role is to faithfully represent the outside world, pixel by pixel, frequency by frequency, and pass on this information to higher sensory areas that interpret this information at higher structure and within context. Such a view no longer seems to be strictly accurate, as evidence of both top down modulation of sensory responses (Niell and Stryker, 2010; Keller, Bonhoeffer, and Hübener, 2012; Shuler and Bear, 2006) and purely non-sensory responses (Stringer et al., 2019; Musall et al., 2019; Steinmetz et al., 2019; Guitchounts et al., 2020) emerge in V1. However, though we now know there is strong modulation of sensory cortical activity by non-sensory events, such as movement, outcome, or expectation, there have been conflicting reports about the specificity of the information represented about these variables in early sensory cortex (Steinmetz et al., 2019; Guitchounts et al., 2020).

Here, I interrogated neuronal responses in V1 during a visual discrimination task to understand what level and specificity of information is available to this early stage in the visual pathway. Further, I wanted to understand the organization of these non-sensory representations, relative to the classical visual tuning of individual cells. Finally, I wanted to exploit the previously described behavioral findings, as well as

the stimulus-independent variant of the task, to understand whether and how the sensory demands of the current environment shape information processing in single neurons and across the population. By testing trained animals' responses to atypical stimuli, I was able to expose the underlying strategy of the animal, whereby at baseline, only the information in the lower visual field drives choice behavior, regardless of the information available in the upper visual field. This suggests that there may be an intrinsic "attentional" imbalance between processing of the two information streams. Whether this influences either sensory or nonsensory representations in primary visual cortex is unknown; many attention studies in primates have found that attentional modulation primarily affects higher visual areas downstream of V1. However, the intrinsic (rather than experimenter-cued) 'attention' phenomenon observed here may offer a different view.

## 4.2 Sensory and non-sensory representations in V1 during the cloud of dots task

I used tetrode drives to record in V1 of freely moving animals performing the cloud of dots discrimination task (Figure 4.1A). Using semi-automated clustering (Kilosort, Pachitariu et al., 2016, Methods) with manual curation, I identified 433 putative single units over 48 sessions across 4 animals on the basis of spike amplitude, spike shape, ISI distribution, and auto- and cross-correlograms. The mean activity patterns of individual neurons across trials showed peak activity at timepoints spanning the trial, with over half of recorded units showing maximum activity during the choice and outcome periods (Figure 4.1B). I first sought to characterize the structure of task representations within single neurons. For classification of single neuron response profiles, I used only well-isolated units with ISI violations below 1%. This left 345 cells for the following analyses. I used a selectivity index to identify selective cells on the basis of the normalized difference in mean firing between two conditions A and B:

$$si = \frac{FR_A - FR_B}{FR_A + FR_B}$$

I used this measure to quantify the selectivity for easy upper versus easy lower stimulus during the stimulus epoch; left versus right choice during the choice epoch; and reward versus punishment during the first 500ms of the outcome epoch. I compared the value of the selectivity index for a given variable for each cell to the shuffle control selectivity distribution for that cell, and identified as selective any cell that had a selectivity index that exceeded the 95% confidence interval of the shuffle distribution for the same variable. The shuffle control selectivity distribution was built by shuffling the condition labels relative to the firing rates, keeping the same number of datapoints for each condition as in the unshuffled data, and calculating the selectivity index by the above equation. This was done 1000 times for each cell to generate a distribution of selectivity indices. I then calculated the overlap in the cells selective for different features to look for any significant enrichment of selectivity for a specific combination of features. The selectivity profiles for the above three features is shown in Figure 4.1C, with the inset showing the expected probabilities for each selectivity profile given the proportions of cells selective for each individual feature under the assumption that selectivity for each feature is independent. The proportions of cells selective for each combination of features is similar to that expected from the joint probabilities of independent events. For example, the number of selective cells found in the data for the stimulus+choice+outcome selectivity profile falls within the 95% confidence interval of the distribution of expected cell proportions, computed using bootstrap sampling (10000 repeats). All other categories had a smaller magnitude of difference between the expected and measured category proportions.

To understand the information encoded in the population, I first identified cells that were selective for the spatial distribution of the cloud of dots stimulus. Two oppositely tuned cells are shown in Figure 4.2A-B. I used a stimulus selectivity index to measure the degree of tuning of a neuron's response during the stimulus period to 95% upper versus 95% lower stimuli:

$$si = \frac{FR_{upper} - FR_{lower}}{FR_{upper} + FR_{lower}}$$

As described above, a cell was classified as stimulus selective when its selectivity index si exceeded the 95% confidence interval of the shuffle control (Methods). This analysis was restricted to cells that

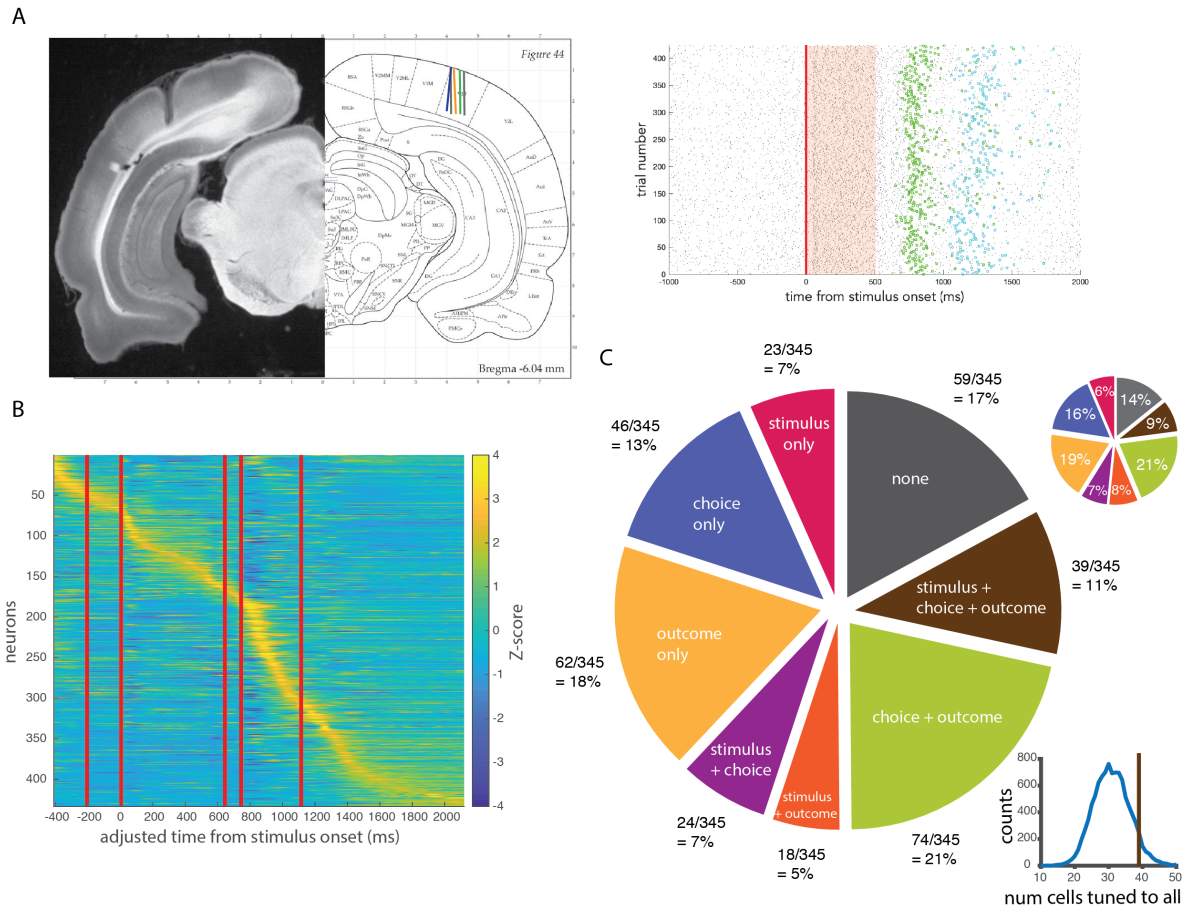


FIGURE 4.1: **Response profiles in V1 of behaving animals:** A. Left: Recording sites over animals. Colored lines indicate recording sites of the corresponding animal whose behavior was shown in Chapter 2. Grey lines indicate recording sites in two animals trained on passive viewing only. Right: Single neuron raster plot aligned to stimulus onset (red vertical line) over trials. Green and blue ticks indicate movement and reward onset, respectively. B. Mean trial-aligned z-scored activity for all recorded units in the cloud of dots task (left, N=4,n=433). C. Proportions of well-isolated single units selective for stimulus, choice, or outcome in their respective epochs, and combinations of these features. Top inset: expected mean proportions given independent sampling of the three features. Bottom inset: Distribution of expected number of cells (blue) selective for stimulus, choice, and outcome computed using bootstrap sampling over 10000 repeats. Measured number of cells (n=39, brown) with this selectivity profile falls within the 95% confidence interval of this distribution (p = 0.13).

fired on average more than 1 spike/s during the stimulus epoch (n = 249). By this metric, I identified two populations of cells, one of which preferred when most dots were in the lower visual field (lower preferring cells, n=56), while the other preferred when most dots were in the upper visual field (upper preferring cells, n=48). More than 50% of well-isolated single units with greater than 1 spike/s firing rate in the stimulus period were assigned as nonselective (n = 145, Figure 4.2D). Mean activity trajectories were similarly heterogeneous across the two stimulus selective subpopulations, over task epochs (Figure 4.2C). Mean firing rate distributions did not differ significantly between the two stimulus selective subpopulations or with the nonselective cells. These classifications corresponded well with a second selectivity measure based on receiver operating characteristic (ROC) analysis, which calculates the probability that the stimulus identity (easy trials only) can be decoded from single trial firing rates (Figure 4.3A).

While almost half of all cells were significantly stimulus selective, all cells performed less accurately than the animal in terms of discriminating between stimuli. ROC analysis was used to estimate the sensitivity of single cells to different stimulus conditions (Methods, adapted from Britten et al 1992). Briefly, I matched stimulus conditions across the 50%-50% decision boundary, and used the area under the ROC curve to estimate the discriminability of a single cell's stimulus epoch firing rate between matched stimulus conditions. Doing this for 95%-5%, 75%-25%, and 55%-45% comparisons produced neurometric curves aligned to the comparison strength of the stimulus. Two example neurometric curves corresponding to the best lower-selective (blue) and upper-selective (orange) cell are shown in Figure 4.3B. In both cases, the neurometric curve is shallower than the psychometric curve describing the behavior of the animal. Indeed, all identified spatially-selective cells in this dataset had shallower neurometric curves than the psychometric curves on that session (Figure 4.3). In this dataset, V1 single neurons are uniformly less sensitive to the spatial distribution of the stimulus than the animal reported through behavior.

In addition to stimulus selectivity, single neurons also showed selectivity to choice side and trial outcome, often in combination (Figure 4.4A-C). Individual cells either showed sustained or consistent tuning across task epochs (e.g. Figure 4.4A top: example cell preferring right choice prior to, during,

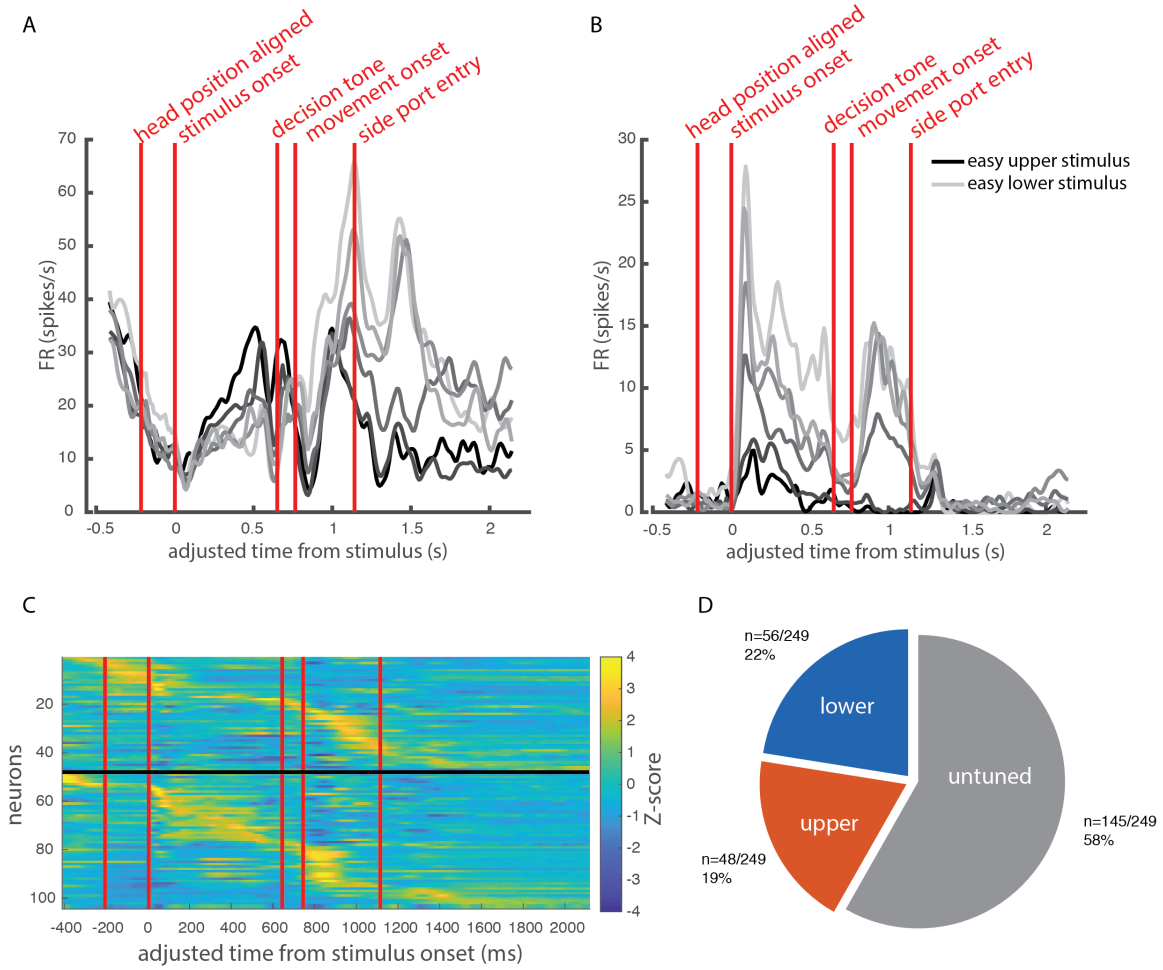


FIGURE 4.2: **Representations of visual space in V1:** A,B. Example putative single neurons showing tuning to spatial distribution of flickering "cloud of dots" stimulus. PSTHs were "stretched" on a per-trial basis to align to all displayed task events (red vertical lines, labels as in left panel, Methods). A. Neuron preferring stimuli with more dots in the upper half ("upper preferring"). B. Neuron preferring stimuli with more dots in the lower half ("lower preferring"). C. Identified subpopulations of stimulus selective cells show similar mean activity patterns over the course of the trial. Vertical red lines denote event times as in panel B. Cells above the black line are upper-preferring, while cells below the black line are lower-preferring. D. Upper-preferring (orange) and lower-preferring (blue) cells are represented in the recorded population to roughly equal numbers. Over half of putative single neurons with activity during stimulus presentation (n=145/249) did not show significant stimulus tuning.

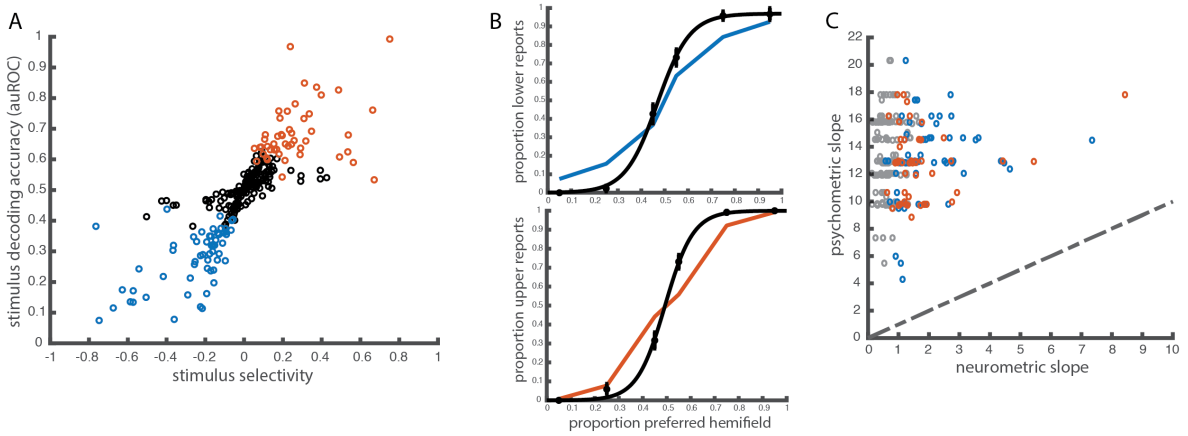


FIGURE 4.3: **Single V1 neurons are less sensitive than the animal's behavior:** A. Comparison of selectivity index metric to auROC. B. Best single neurons in each subpopulation (blue = lower preferring, orange = upper preferring) are less accurate than animal behavior (black) on the same session. Neurometric is calculated using ROC analysis, per Britten et al 1992 (Methods). C. Across all stimulus-selective cells (blue = lower preferring, orange = upper preferring), slope of the neurometric is much shallower than the slope of the associated psychometric. Untuned neurons (grey points) have very low neurometric slopes, as expected. Slopes are estimated from a logistic function fit to either the neurometric or behavioral data. Dashed line indicates the unity line. This trend is consistent when the neurometric slope is calculated as the linear slope between individual points.

and after termination of choice report) or more transient selectivity in different task epochs (e.g. Figure 4.4A bottom: example cell with transient, epoch-specific side preference between choice and outcome epochs). Outcome-selective responses could either be transient or sustained until the last exit from the side port (Figure 4.4B-C), and typically differentiated well between rewarded and unrewarded (including both punished and missed reward) trials. In addition, most outcome-selective cells were differentially selective based on choice side (Figure 4.4C). Across the population, almost 2/3 of all well-isolated cells were choice selective, and a similar proportion were outcome-selective (Figure 4.4D).

It has been postulated that movement responses in V1 may reflect either sensory reafference or corollary discharge from motor command neurons. It has also been reported that movement-related responses in V1, particularly those that emerge immediately prior to the movement onset, which cannot be due to sensory reafference, carry no information about the impending choice (Steinmetz et al., 2019). However, a subset of the choice selective neurons identified on this task (14%, n = 26/183) began to show choice direction selectivity prior to movement onset. In some cases, this tuning was in the same direction as tuning in the movement epoch (Figure 4.5A, top), whereas other cells showed more complex dynamics, with opposing preferred choice directions between the pre-movement and movement epochs (Figure 4.5A, bottom). In both of these cells, the tuning direction in the pre-movement epoch is discontinuous with the tuning in the stimulus epoch, and is therefore unlikely to reflect continued activity based on stimulus preferences.

Indeed, there was not always a clear relationship between stimulus tuning and choice tuning in cells that exhibited both. I found that when considering cells with firing of greater than 1 spike/s in both stimulus and choice epochs, about a quarter of these cells (n = 63/239) were significantly tuned during both epochs. However, the tuning was not always in the same direction. Two example cells shown in Figure 4.5B illustrate that some cells preferred the opposing choice direction during the choice epoch rather than the direction that corresponded to their preferred stimulus class. These PSTHs were restricted to only correctly reported trials, where there is a one-to-one relationship between stimulus classification and choice.

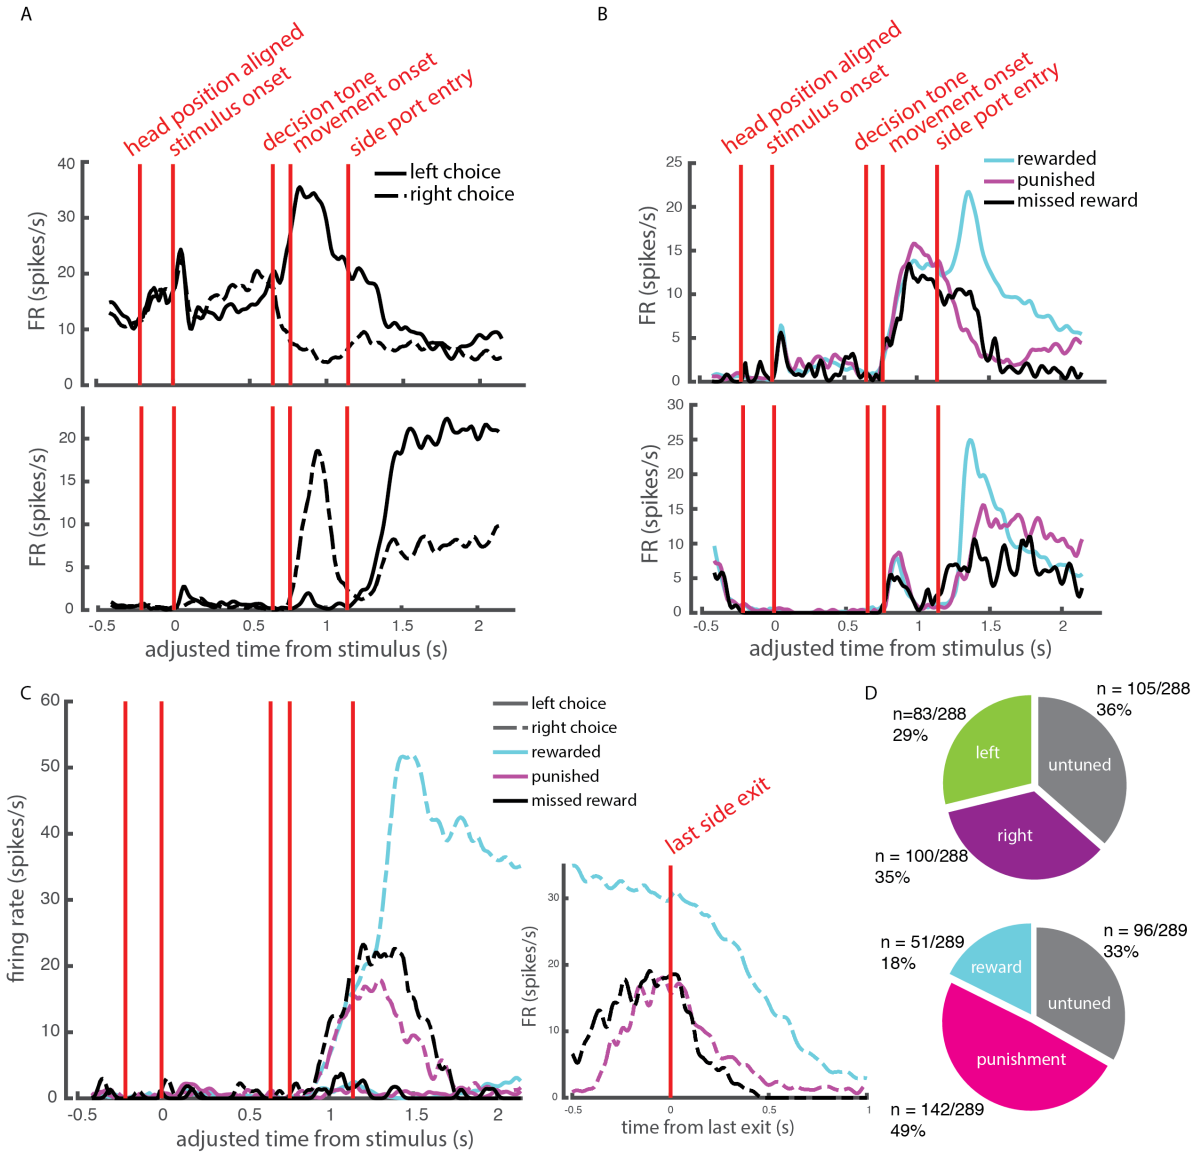
A
head position aligned
stimulus onset
decision tone
movement onset
side port entry
left choice
right choice
FR (spikes/s)
adjusted time from stimulus (s)
B
rewarded
punished
missed reward
C
firing rate (spikes/s)
last side exit
time from last exit (s)
D
n=83/288
29%
left
untuned
n = 105/288
36%
right
n = 100/288
35%
reward
n = 51/289
18%
punishment
n = 96/289
33%
n = 142/289
49%

The fact that choice tuning in single cells can be independent of and sometimes opposite to tuning for the associated stimulus, for example, seems consistent with the idea that single V1 neurons reflect activity during movement onset that is merely the corollary discharge from motor effector neurons, carrying pure motor information that is irrelevant to the immediate local computations. I hypothesized that if this were the case, then all movements of a similar type, such as between-port movements, should elicit similar responses and reflect similar tuning within single neurons. In the course of a single trial, an animal makes two between-port movements: one to report the choice (center-to-side) and one to return to re-initiate the next trial (side-to-center). Across many trials, both of these movements will cover both directions relatively equally. I expected that if V1 neurons were faithfully encoding movement direction, that they would fire similarly for e.g. a rightward choice movement (center port to right port) as for an initiation movement leaving the left port (left port to center port). Instead, I found cells that had opposing tuning preferences (e.g. Figure 4.5C, example cell 1) or differing degrees of tuning (e.g. Figure 4.5C, example cell 2) across the two movements. For this analysis, I restricted initiation movements to those that were completed in $< 0.5$s between side port exit and center port entry, to remove any indirect movements that may include exploration of the behavior box, pauses, grooming, and the like. When I compared the selectivity indices between these two movement epochs, again calculated on only direct port-to-port movements, there was no significant correlation across the population between tuning direction and magnitude across these two epochs (Figure 4.5D).

From the example single neuron PSTHs in Figure 4.5A-C, it appears that various combinations of tuning preferences across the different behavioral epochs emerge at the level of single cells. I quantified this relationship by comparing the selectivity indices of individual cells for stimulus compared to those for choice probability (i.e. selectivity during stimulus period firing for future choice), choice side,

FIGURE 4.4 *(preceding page)*: **Nonsensory representations in V1 of behaving animals:** A. Two example cells showing choice direction tuning. Vertical red lines denote the event timepoints shown above. B. Two example neurons showing outcome selectivity. C. Outcome responses are often choice-side dependent. Inset shows the same neuron's firing aligned to the last side port exit, demonstrating that the increased activity on rewarded trials exceeds what is due to remaining at the side port for a longer duration. D. Proportions of putative single neurons with activity during the relevant epochs above 1 spike/s showing tuning to choice side (top) and reward delivery (bottom).

and reward. In each case, I found low and/or insignificant correlations between these different sets of selectivity indices (Figure 4.5E). By the two sample Kolmogorov-Smirnov test, there is no significant difference in the distribution of choice selectivity between upper- and lower-preferring subpopulations ($p = 0.876$). However, there is a significant difference in the distributions of choice probability, defined here as the deviation from chance of the area under the receiver operating characteristic curve calculated from stimulus period firing for upcoming choice ($p = 0.0232$), and the distribution of reward selectivity ($p = 0.0007$) across these two stimulus selective subpopulations.

## 4.3 Visual and non-visual responses in a stimulus-independent task

I also recorded V1 activity in animals performing the stimulus-independent task, where visual stimulus was uninformative for choice. I recorded a total of 356 units across 3 animals. Of these, 255 units were recorded in 2 animals presented with the same overall stimulus density as in the discrimination task, but where individual dots were randomly distributed across the screen, whereas 101 units were recorded in an animal presented with a sparser stimulus. I focus on the first dataset for the following analysis. Of these 255 units, I identified 177 well-isolated cells with fewer than 1% refractory period violations; I will use these for all further single cell response characterizations.

I first characterized the dynamics of the responses in V1 across all recorded units. In both the cloud of dots and stimulus-independent tasks, fewer than a third of recorded units showed their greatest response during the stimulus period, while many units showed maximum responses during the choice and outcome epochs (Figure 4.6A-B). A few of these cells (e.g. those at the very top of the plot) showed responses during both movement into the center port (prior to the first timepoint of head position alignment) and movement out of the center port (time period between the last two timepoints), but most movement selective cells showed much greater responses in one of these periods than the other.

When the same selectivity index analysis described above was applied to the firing rates of well-isolated cells during presentation of the easy-upper and easy-lower stimuli from the discrimination task, approximately 40 % of cells were identified as being stimulus selective (Figure 4.6C). This dataset

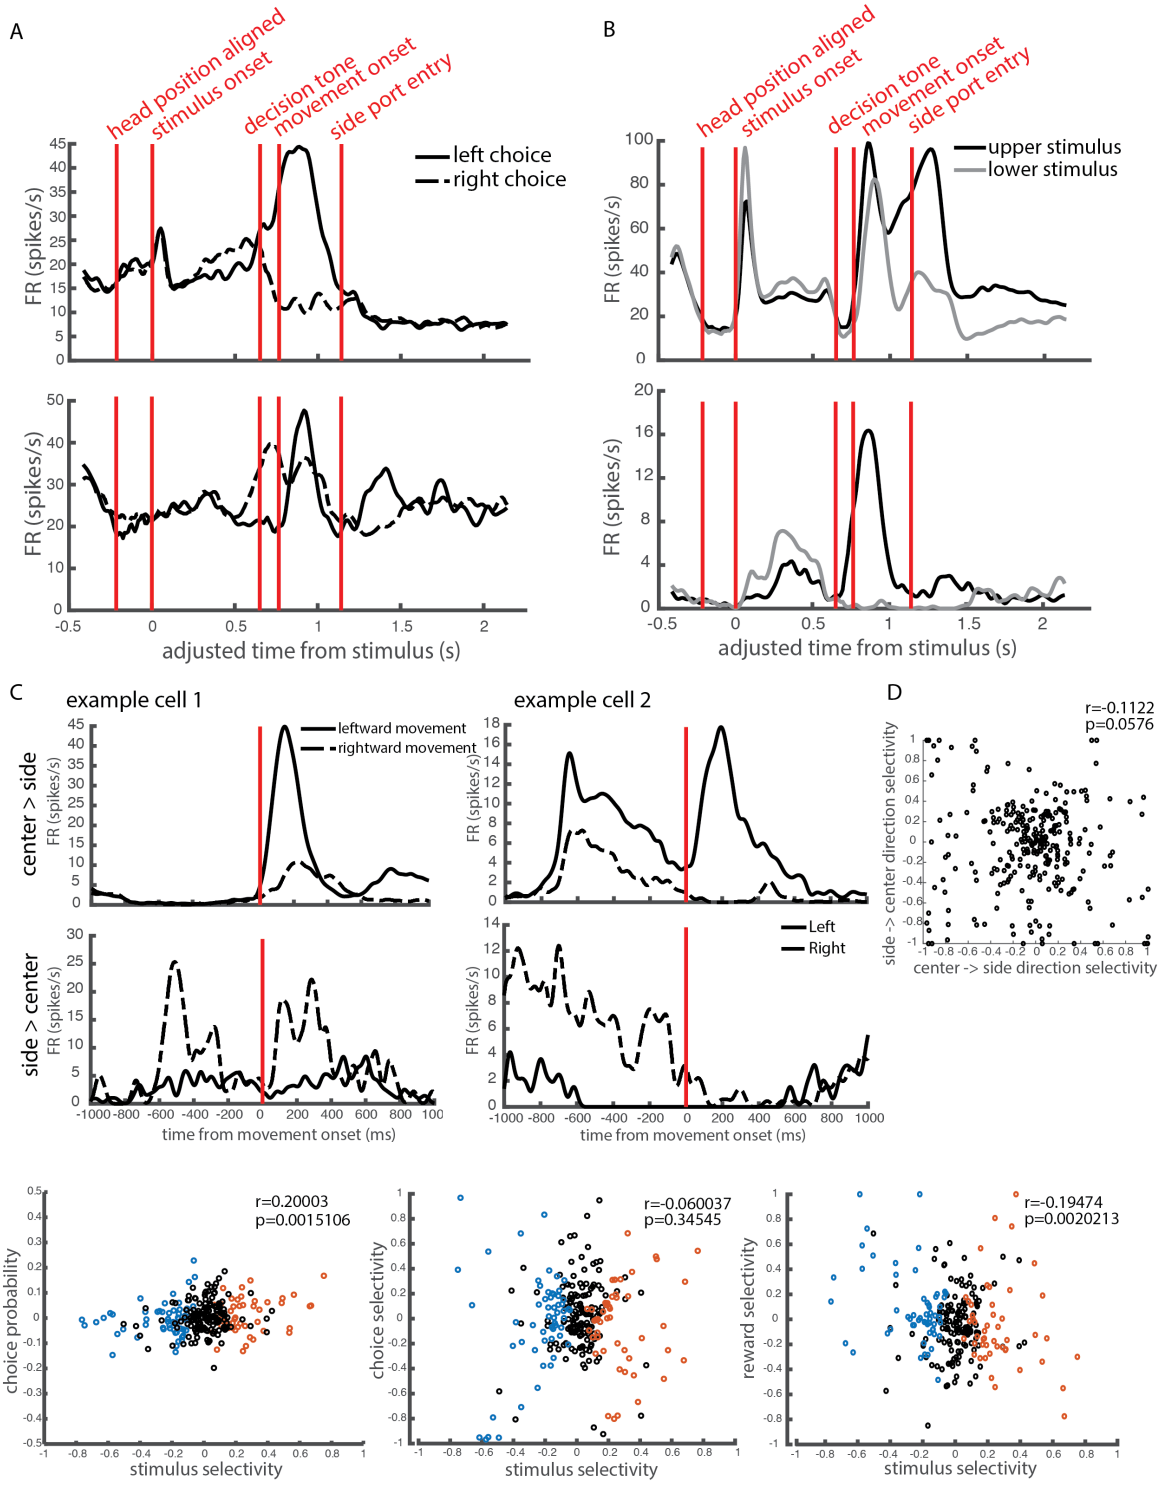
A
head position aligned
stimulus onset
decision tone
movement onset
side port entry
left choice
right choice
FR (spikes/s)
adjusted time from stimulus (s)
B
upper stimulus
lower stimulus
C
example cell 1
example cell 2
leftward movement
rightward movement
center > side
side > center
Left
Right
time from movement onset (ms)
D
r=-0.1122
p=0.0576
side -> center direction selectivity
center -> side direction selectivity
E
r=0.20003
p=0.0015106
choice probability
stimulus selectivity
r=-0.060037
p=0.34545
choice selectivity
r=-0.19474
p=0.0020213
reward selectivity

was dominated by lower-preferrring cells; this is likely due to slight variation in the location of the recording site in these two animals relative to their retinotopic maps. Similarly to cells recorded on the discrimination task, there was weak or insignificant correlation between stimulus selectivity and choice or outcome selectivity (Figure 4.6D). Here, there are fewer cells for which I could calculate the choice selectivity, because in this task animals tend to visit predominantly one choice side over the other, and so only cells that were recorded on sessions including trials of both choices qualified for this analysis.

I wanted to compare features of the movement-related responses between the "choice" epoch, and the trial initiation epoch. Similar to above, I used only trial initiation movements that were completed in less than 0.5s, a comparable timescale to the majority of choice movements. Because choice selectivity classifications were sparse, and expected to be slightly noisy, I compared the amplitude and latency of responses between these two epochs across all well-isolated single units (Figure 4.6E,F). A majority of cells had a higher peak response during the choice movement than during initiation movement (n = 113/177, 64%), while the remaining cells fired more strongly during the initiation movement. Comparing the latency from movement onset at which this peak response occurred showed no correlation within cells between latency to respond during the choice epoch and during the initiation epoch. The same comparison is shown for cells recorded on the discrimination task (Figure 4.6G,H). Here, roughly equal proportions of cells fired more strongly in the choice epoch versus in the initiation epoch (n = 150/288, 52% with higher peak firing in the choice epoch). This was similar when only considering movement

FIGURE 4.5 *(preceding page)*: **Movement responses in V1 differ across task epochs:** A. Two example cells showing choice direction tuning beginning prior to movement onset. Cells can be tuned in the same direction in the movement preparation epoch (between decision tone and movement onset) as in the movement epoch (top) or in the opposite direction (bottom). Vertical red lines denote the event timepoints shown above. B. Two example neurons showing inverted tuning during movement compared to stimulus epoch. Top: A cell preferring lower stimuli during stimulus presentation, but the upper associated choice during the movement and outcome epochs. C. Two example neurons exhibiting choice direction tuning that do not exhibit movement direction tuning in the same direction (example cell 1) or at all (example cell 2) during the return movement from the side to center port. D. Comparison of between-port movement direction preference during the choice epoch versus prior to trial initiation. E. Representations of nonsensory parameters such as choice (either choice probability during the stimulus period or choice selectivity during movement) and reward show weak or no correlation with representations of visual stimuli (r and p values shown for each plot). Blue points denote lower-preferring cells, while orange points denote upper-preferring cells. Nonselective cells are shown in black.

direction-selective cells (n = 98/183, 54%). Similarly to the passive task, there was no correlation between the latency to peak response across these two epochs.

## 4.4 Variance of visually evoked responses depends on task demands

Trial-to-trial variability has been used as a measure of the reliability of a cell's evoked response to a given stimulus, especially in the attention literature. The implication is that when cells fire more consistently over repeats of the same condition, this facilitates more reliable readout of their activity by downstream circuits. As a result, measures such as the Fano factor, which measures the variance of a cell's activity in a given epoch relative to its mean firing rate, have been studied extensively in the attention literature (McAdams and Maunsell 1999, Mitchell et al 2007). However, most findings of decreased Fano factor have been found in V4 of the macaque, several stages past primary visual cortex, in the context of cued attention tasks. These changes often appear in combination with increased firing rates and changes in the correlation structure between neurons. The Fano factor is defined as:

$$fano = \frac{var_{spikecount}}{mean_{spikecount}}$$

I wondered here if the Fano factor would vary across stimulus-selective subpopulations in the discrimination task, given their differential behavioral relevance. Indeed I found that despite the fact that there was no significant difference in the distribution of mean evoked firing rates during the stimulus epoch in upper-preferring versus lower-preferring cells during presentation of their preferred stimulus (p = 0.581), the distribution of Fano factors in lower-preferring cells was significantly lower during the same task epoch (p = 0.0015, Figure 4.7A-B). In the stimulus-independent viewing task, in contrast, the lower preferring cells did not have lower Fano factors than either the untuned or upper-selective cells (Figure 4.7E). In fact, in these data, the upper-preferring cells had lower Fano factors in the viewing task (p=0.000028); however, as noted before, only 10% of well-isolated cells (n = 17/163) were identified as upper-preferring, so outliers may have an exaggerated effect on these data.

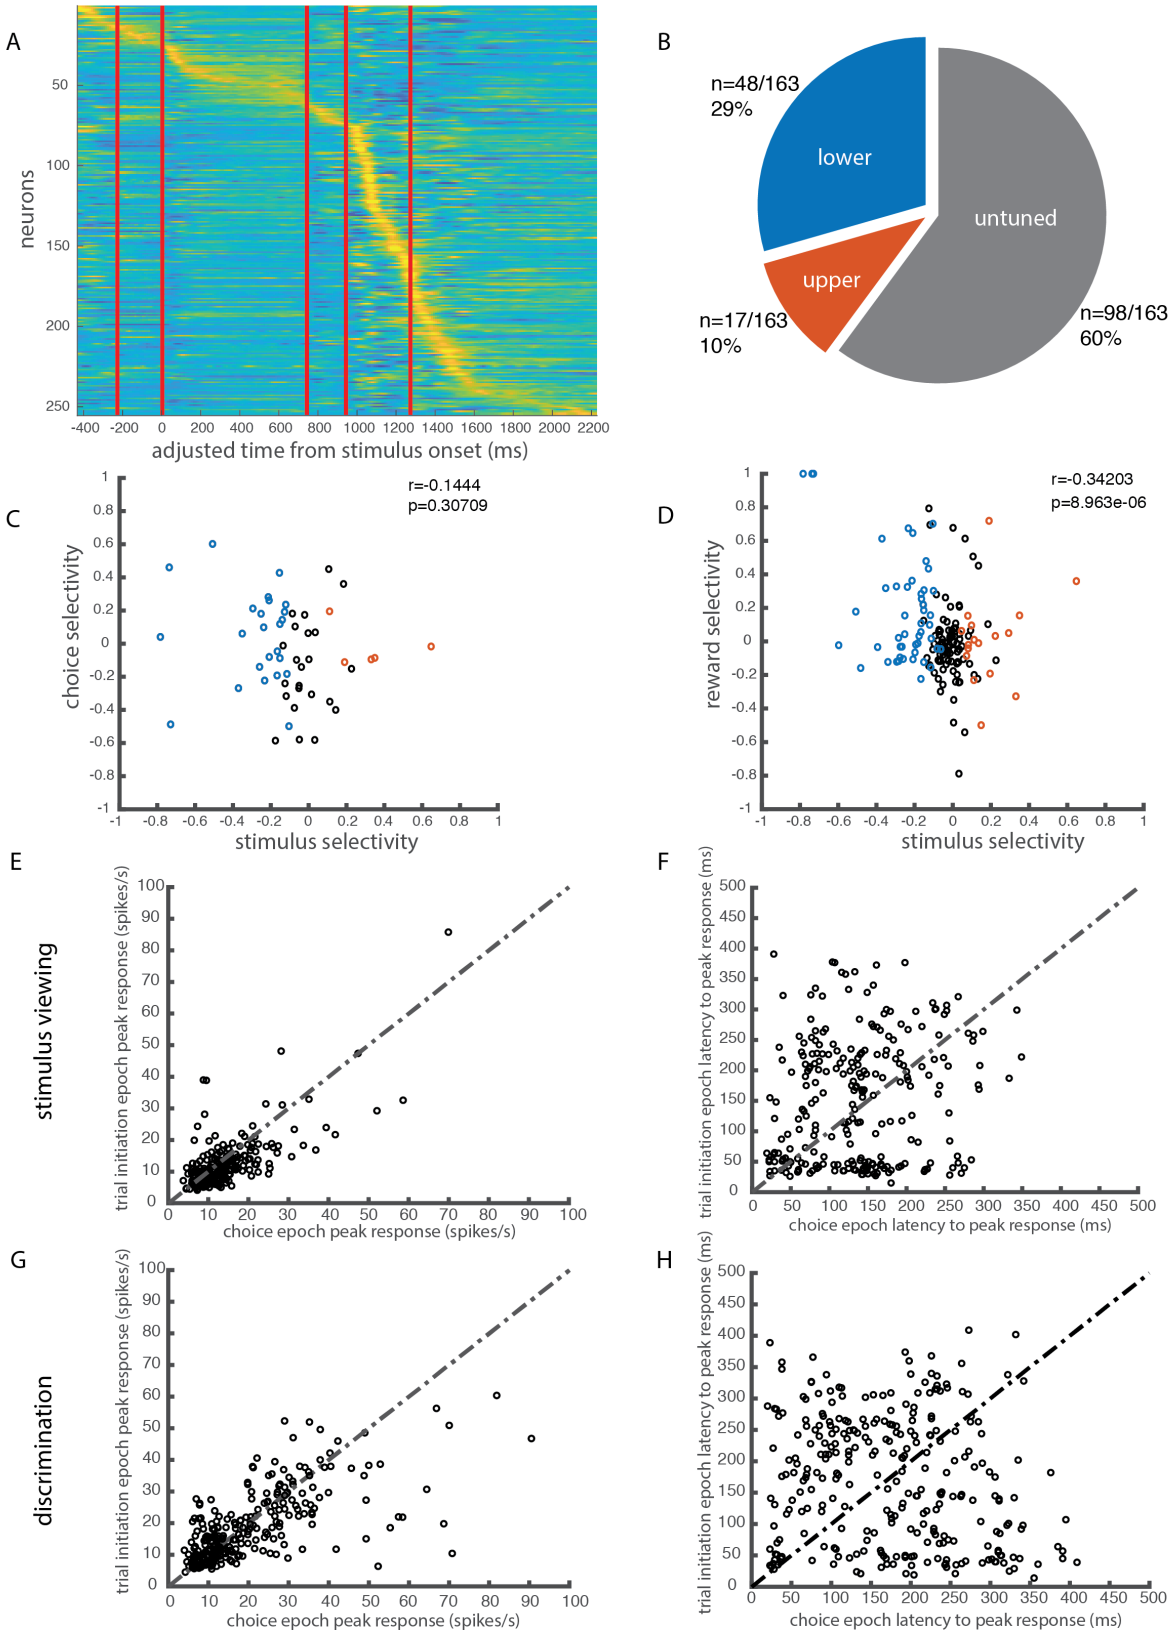
A
neurons
adjusted time from stimulus onset (ms)
B
n=48/163
29%
lower
untuned
upper
n=17/163
10%
n=98/163
60%
C
r=-0.1444
p=0.30709
choice selectivity
stimulus selectivity
D
r=-0.34203
p=8.963e-06
reward selectivity
stimulus selectivity
E
stimulus viewing
trial initiation epoch peak response (spikes/s)
choice epoch peak response (spikes/s)
F
trial initiation epoch latency to peak response (ms)
choice epoch latency to peak response (ms)
G
discrimination
trial initiation epoch peak response (spikes/s)
choice epoch peak response (spikes/s)
H
trial initiation epoch latency to peak response (ms)
choice epoch latency to peak response (ms)

I was curious as to whether the difference in Fano factor was specific to the stimulus epoch, or whether it reflected sustained differences in firing across the full trial. To test for this, I plotted the timecourse of the mean Fano factor for each of the stimulus-selective and the non-selective subpopulations, calculated in sliding bins of 50ms and aligned to stimulus onset (Methods). Here, I saw a clear difference in how the Fano factor evolved over the course of the stimulus period. Specifically, while the mean Fano factor across lower-preferring cells saw a stimulus onset-aligned decrease (or "quenching") relative to baseline that was sustained over the duration of the preferred stimulus, the upper-preferring and nonselective subpopulations did not decrease their variability at the onset of their preferred stimulus (Figure 4.7C). The difference in the mean Fano factor between the lower-preferring and upper-preferring subpopulations was significant for the majority of the stimulus period (unpaired t-test, $p<0.05$). Interestingly, there was a gradual increase in the Fano factors within the upper-preferring subpopulation over the course of the stimulus period, that decreased back to baseline following stimulus offset. This pattern was not observed in the non-selective subpopulation. Previously, increases in Fano factor have been linked to accumulation processes, where independent noise accumulates as signals are integrated over time. It will be interesting to explore the robustness of this effect over a larger dataset, and to understand how this relates to single cell firing patterns within the subpopulation. Repeating this analysis with a longer time window (100ms or 200ms) led to smoother trajectories but similar overall patterns between the three subpopulations. Similarly, repeating the analysis using the mean-matched Fano factor, which takes into account and subsamples to remove the effect of firing rate differences, produced similar patterns to the raw Fano factor analyses shown here.

FIGURE 4.6 *(preceding page)*: **V1 mean selectivity patterns and distribution of task representations during a stimulus-independent viewing task:** A. Mean trial-aligned z-scored activity for all recorded units in stimulus-independent viewing task (right, N=2,n=255). Event times are denoted by red bars, and follow the same order as in figure 3.1. B. Proportions of putative single neurons recorded under viewing conditions showing selectivity to easy cloud of dots stimuli. Few upper-preferring cells were identified, likely due to slight differences in position of the recording site along the anteroposterior axis. C,D. Comparison of distributions of nonsensory representations such as choice side (C) and reward delivery (D) relative to stimulus tuning. E. Comparison of peak response amplitude during choice epoch (x-axis) and trial initiation epoch (y-axis) during the passive stimulus viewing task. F. Similar to E, but comparing latency to peak response across the two movement epochs in the passive stimulus viewing task. G, H. Same as E,F but for discrimination task.

In contrast, stimulus-tuned subpopulations on the passive task did not show the same patterns when compared on the same analysis (Figure 4.6F). Instead, all subpopulations showed a slight decrease in Fano factor for the duration of the stimulus period. The similarity between the subpopulations on this analysis was somewhat surprising given the much lower Fano factor for the upper selective cells during the full stimulus period. The lower Fano factors across the upper-preferring subpopulation became apparent as I increased time windows to 200ms in this analysis (Figure 4.7G), as did the stimulus onset-evoked dip in variability across all subpopulations (Figure 4.7F-G). This change based on the size of the sliding window was consistent with previous reports that Fano factor estimations are dependent on the length of the time window used to calculate them (Cohen and Kohn, 2011). However, it is notable that increasing the time window from 50ms to 200ms did not substantially change the effect size between the lower- and upper-preferring subpopulations in the discrimination task (Figure 4.7F), where the Fano factor across the two populations remains significantly different throughout most of the stimulus period. In the passive task, however, though the upper selective cells have lower Fano factors across the population than the lower selective cells, this emerges before stimulus onset and is therefore more likely to reflect intrinsic differences in the cells sampled (e.g. firing rate outliers or cell type, Figure 4.7G).

## 4.5 Population decoding improves with population size, carries more info about past choice than upcoming choice

I wondered if readout of features of the stimulus and task variables improved at the level of the population relative to the component individual cells, and what role non-selective cells played in this improvement relative to selective cells. In this dataset, I recorded 48 sessions with simultaneous recordings of from 1 to 36 cells. I restricted this analysis to sessions with 4 or more simultaneously recorded cells. This left 26 sessions for the following analysis. I then attempted to decode either stimulus, choice, or outcome from the vector of firing rates across simultaneously recorded cells on a given trial. For this analysis, I used a linear classifier implemented by the Matlab fitclinear function, with L1 regularization and 10-fold

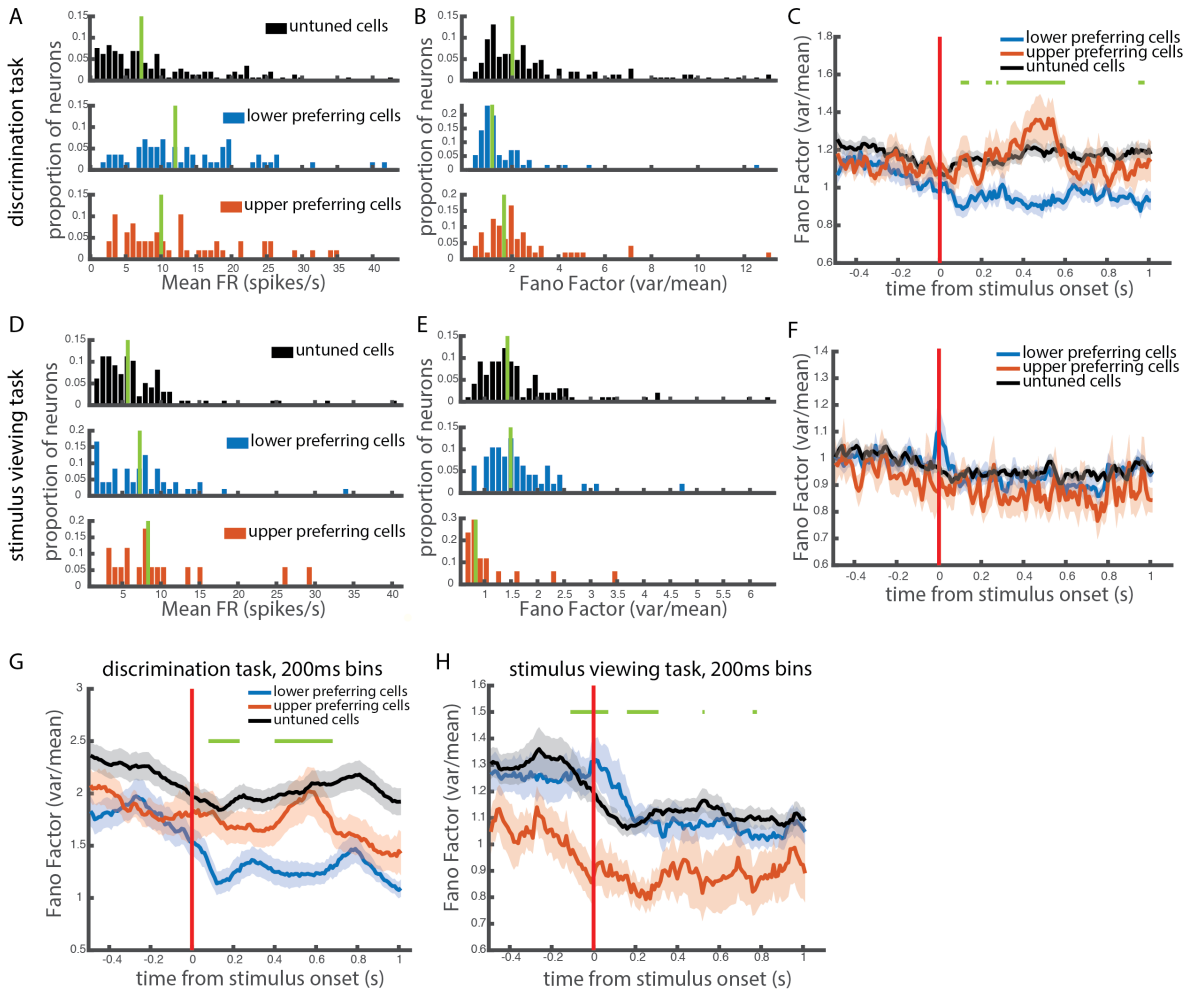


FIGURE 4.7: **Trial-to-trial variability of stimulus representations depends on behavioral demands:** A,D. Mean firing rates during preferred stimulus presentation for tuned subpopulations in the cloud of dots discrimination task (A) and passive viewing task (D). Green vertical lines denote the median across the population. B,E. Fano factor during preferred stimulus presentation for tuned subpopulations in the discrimination task (B) and passive viewing task (E). C,F. Fano factor as a function of time from onset of preferred stimulus, for stimulus-tuned and untuned subpopulations in the cloud of dots discrimination task (C, stimulus duration 0.5s) and passive viewing task (F, stimulus duration 0.7s), calculated in 50ms sliding bins. Green bars denote timepoints where the distribution of fano factors between upper-preferring and lower-preferring subpopulations are significantly different (unpaired t-test, $p<0.05$). Shaded regions indicate one standard error of the mean. G,H. Same as C and F respectively, but using 200ms sliding bins. Note: timecourse is plotted aligned to the final timepoint of each bin (i.e. point at time 0 refers to the bin from 200ms before stimulus onset to time of stimulus onset) leading to a slight lag in the apparent onset and offset of the effect between G,H and C,F.

crossvalidation.

I found that for both stimulus and future choice (choice probability) decoding from stimulus epoch firing rates, classification accuracy on the test set (Methods) improved with population size. However, there was a large spread across sessions with a given population size (Figure 4.8A). Across all sessions, the decoding performance of the full population exceeded the performance of its best component cell (Figure 4.8B). This pattern was the same when I tried to decode upcoming choice from stimulus period activity, though the accuracy for choice probability was in general lower than for stimulus (Figure 4.8C-D). When I applied the same analysis to decode previous trial choice from stimulus period activity, I observed slightly higher accuracy across all sessions, which also scaled with population size (Figure 4.8E). Finally, and unsurprisingly, decoding choice direction from movement epoch activity yielded the highest accuracy (Figure 4.8F).

I finally tested for the contribution of stimulus selective and non-selective cells to decoding stimulus identity from stimulus period firing rates. I split the populations of simultaneously recorded cells into selective and nonselective populations on the basis of the individual selectivity indices of cells. This analysis was restricted to sessions including both types of cells. Although in some cases I was able to decode stimulus identity from subpopulations of nonselective cells only, the decoding accuracy was higher when I used only cells that were defined as individually stimulus selective (Figure 4.8G). Finally, I tested the difference in the the decoding accuracy when nonselective cells were removed versus included in the population, and found that their inclusion had a very small effect on the accuracy of the population. This effect was not significantly correlated with the number of nonselective cells added (Figure 4.8H).

## 4.6 Relationship of recording site to stimulus representations

I recorded 199 cells in 1 animal in an area more anterior to the usual recording site, bridging the very anterior part of V1, anterior V2, and potentially sampling part of posterior parietal cortex. I compared the stimulus representations in this more anterior recording site to those found in the target V1 recording site.

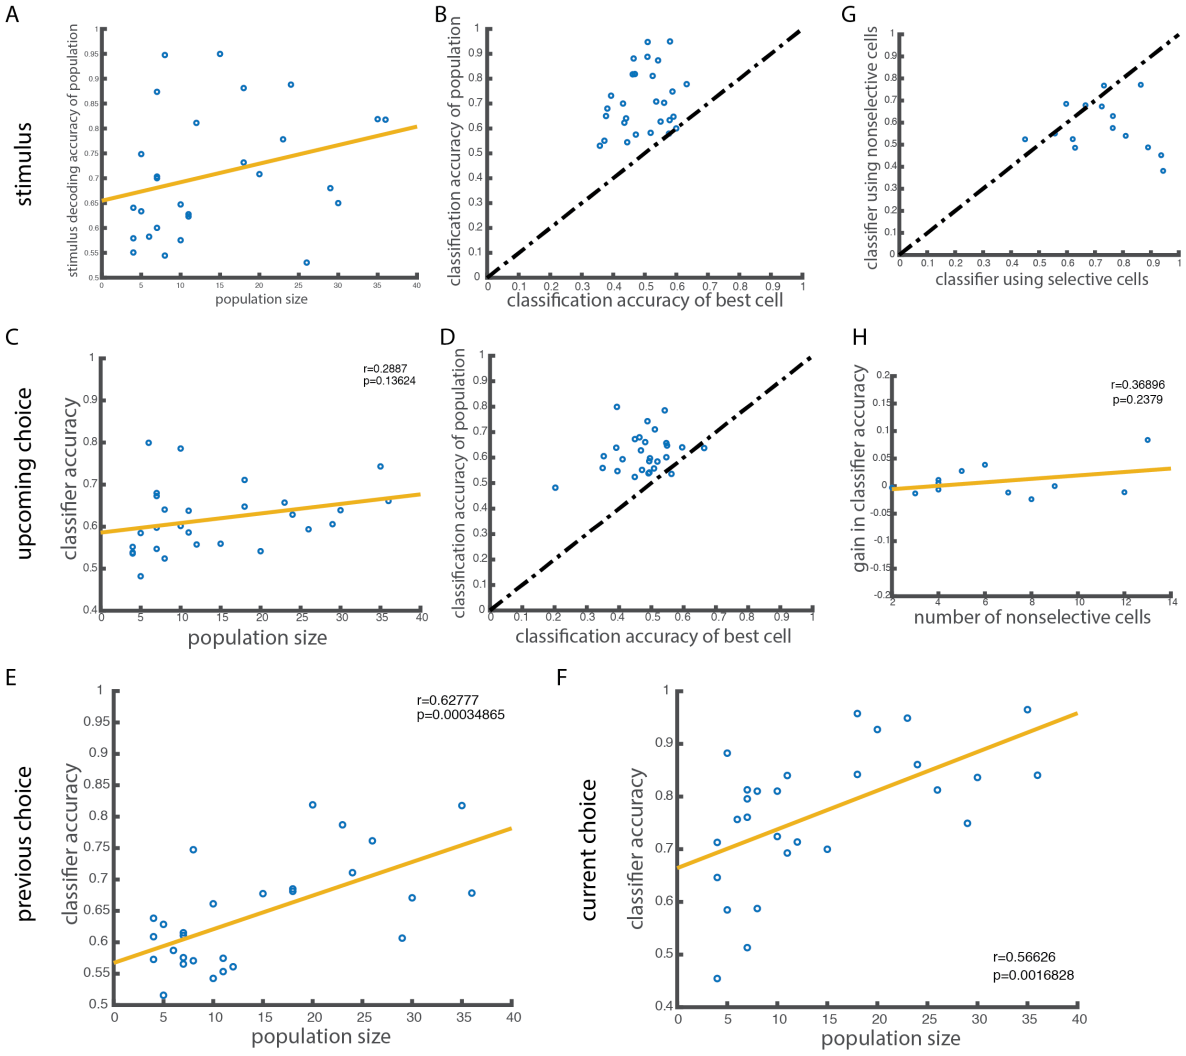
A
stimulus
stimulus decoding accuracy of population
population size
B
classification accuracy of population
classification accuracy of best cell
G
classifier using nonselective cells
classifier using selective cells
C
upcoming choice
classifier accuracy
r=0.2887
p=0.13624
population size
D
classification accuracy of population
classification accuracy of best cell
H
gain in classifier accuracy
r=0.36896
p=0.2379
number of nonselective cells
E
previous choice
classifier accuracy
r=0.62777
p=0.00034865
population size
F
current choice
classifier accuracy
r=0.56626
p=0.0016828
population size

Mean activity trajectories across the population looked similar to those recorded in V1, and a similar proportion of recorded units showed significant selectivity by the above selectivity index analysis (Figure 4.9A,B). Though both the distributions of selectivity indices and decoding accuracy under ROC analysis were slightly narrower than those recorded in V1 proper, these differences were not significant ($p = 0.111$ and $p = 0.339$ respectively, Figure 4.9C compared to Figure 4.3A).

Stimulus period Fano factors were not different across stimulus-selective subpopulations (Figure 4.9D). Analysis of the timecourse of the Fano factor relative to stimulus onset did not show the stimulus tuning-dependent signatures found in V1 proper, though Fano factor in the lower-preferring subpopulation was still slightly lower just prior to and following stimulus onset than non-selective cells. As this recording site was likely to have bridged one or several areal boundaries, it is possible that the estimation of the Fano factor within subpopulations is contaminated by disparate sources of noise from higher visual cortices compared to V1.

## 4.7 Methods

Tetrode drives were custom-built using Omnetics 36-channel EIBs and custom 3D printed drive skeletons. Each drive contained 8 tetrodes and 1 reference tetrode that travelled together in a single bundle. Subjects were implanted with tetrode drives following adequate performance on both the visual discrimination and the head position requirement, typically after at least 1-2 months of training.

I used the Intan-based OpenEphys recording system to record from animals during behavior. Two of the seven animals reported in this chapter required light anesthesia to enable attachment of the recording

FIGURE 4.8 *(preceding page)*: **Information within populations exceeds information in single cells:** A-B. Population decoding accuracy for stimulus identity, from stimulus period activity across simultaneously recorded cells. A. Decoding accuracy increases with population size, but with a large amount of variability. B. Decoding accuracy of the population exceeded in all cases the decoding accuracy of the best cell on that session. C-D. Same as A-B, but for classification of upcoming choice from stimulus period activity. E,F. Same as A, but for decoding previous trial choice from stimulus period activity (E) and decoding current choice from movement period activity (F). G. Decoding accuracy when using only individually stimulus selective cells exceeds accuracy of populations containing only nonselective cells. H. Addition of nonselective cells to the population only slightly improves decoding accuracy from populations of stimulus selective cells only.

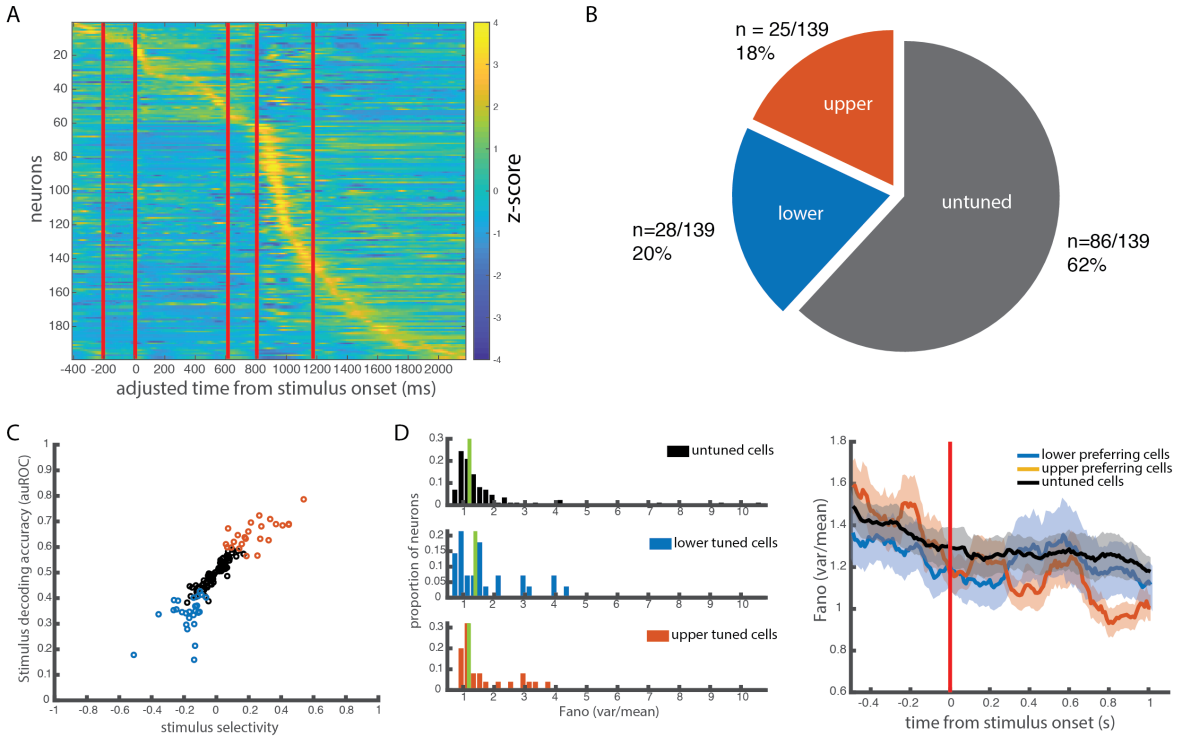


FIGURE 4.9: **Stimulus period response patterns of units at anterior recording site:** A. Z-scored mean activity of all recorded units, sorted by time of peak activity. B. Proportions of stimulus selective cells, defined by selectivity index. C. Correlation between selectivity index and decoding accuracy using ROC analysis. D. Left: Fano factor during preferred stimulus presentation is similar across stimulus-selective subpopulations (all comparisons $p > 0.05$). Green lines indicate median Fano factor across the subpopulation. Right: Fano factor over time, calculated in sliding 50ms bins, is similar during the stimulus period across stimulus-selective subpopulations (paired t-test, $p>0.05$).

tether (1 on the discrimination task and 1 on the viewing task). In these cases, animals were given 15 minutes to fully recover before the task began. During the task, event triggers were continuously sent between the Bpod controlling the behavior state machine and the OpenEphys recording system, via the Open Ephys sync board (Sanworks). Additionally, monitor onsets delivering the stimulus were detected via a custom-built light-sensor module that communicated directly with OpenEphys. After each recording session, I lowered tetrodes by at least 40 μm and no more than 80 μm. Recordings were made until tetrodes reached a depth of 1.5mm, at which point I performed electrolytic lesioning at tetrode tips. Animals were sacrificed and brains were recovered for histology.

I performed semi-automated spike sorting using Kilosort software (Pachitariu et al., 2016) on the raw continuous recording traces. The data was bandpass filtered and the mean across all channels was subtracted from all traces to remove any common noise events. Kilosort was then used to identify and match 128 templates to the raw data. These templates and their associated spike waveforms were manually curated in Phy. During manual curation, I split and merged spikes on the basis of their: amplitudes, waveforms, auto- and crosscorrelograms, firing dynamics over the session, and clustering in feature space. To classify single cell representations, I further curated these units to remove those with refractory period (2ms) violations over 1%. Though visual inspection suggests this spuriously removed several single units with burst-like activity, such cells were not included in the single unit dataset reported above.

All PSTHs here were aligned to the same timecourse, where the length of each variable length task epoch was equalized across all recorded cells and trials. The length of each task epoch was set to the mean epoch length across all trials across all sessions on which units were recorded. Single trial PSTHs were built by calculating the number of spike events in 1ms bins, and smoothed with a Gaussian window of 20ms. Then, uniformly distributed sampling points were calculated based on the event timing on that particular trial, and the single trial PSTH was sampled at those points. This meant that some PSTHs were slightly "stretched" relative to absolute time, and others were "squeezed". However, this allowed us to align activity across trials and across cells to the exact timing of task events such as decision tone, movement onset, and side port entry. Mean PSTHs were then calculated by taking the mean across these

stretched single trial PSTHs for all or a relevant subset of trials.

Mean activity traces across the population were calculated by taking the mean across all stretched single trial PSTHs for each cell, and z-scoring to find the relative activity patterns across the different task epochs. Cells were sorted by the time of maximum activity during the trial.

To find the selectivity of a cell's firing during various task epochs, I calculated a selectivity index using mean firing rates on trial types defined by the task parameter of interest. First, I define stimulus-selective cells as those whose selectivity index exceeds the 95% bounds of a shuffle control distribution. The shuffle control distribution was built by calculating the selectivity index across 1000 shuffles where the label of a given trial (e.g. upper or lower stimulus) was shuffled relative to the single trial firing rates.

ROC analysis was performed using the Matlab `perfcurv` function, using task variable as a binary label, and mean single trial firing rates in a given task epoch as the scores. For comparison with the selectivity index-defined stimulus tuned subpopulations, I applied this analysis to decode the identity of easy stimuli as being either "upper" or "lower" from the stimulus period firing rate. For calculating choice probability, I applied this analysis independently to each of the two difficult trial conditions, and equalized the number of left and right choices within each trial condition by subsampling. Stimulus epoch firing rates were used to decode eventual choice. The choice probability of a cell was defined as the mean of the auROC (area under the curve) for the two trial conditions.

To build the neurometric curve, I applied ROC analysis at each of the 3 difficulty levels presented, and took the area under the curve as the cell's ability to discriminate between the two easy, the two medium, and the two difficult stimuli. These values were mirrored across the 50% point of the decision axis to estimate the full psychometric curve. For comparison of the slopes of the neurometric and associated psychometric curves, I fit a logistic function to the 6 points from the auROC analysis, and a second logistic to infer the psychometric function from the choice behavior, and compared the slope parameter from these two fits.

Fano factor was computed as the variance of the spike count over trials for a given trial epoch divided by the mean spike count in that epoch. For full stimulus period fano factors, this was calculated over the

500ms stimulus epoch. For the Fano factor timecourse analysis, this was calculated on a per cell basis in sliding bins of 50ms every 10ms from 500ms prior to stimulus onset to 1000ms following stimulus onset. I then took the mean of the Fano factor timecourse over cells of interest (e.g. stimulus-selective subpopulations) and calculated the standard error of the mean across those cells' Fano factors. I also repeated this analysis by using a linear regression coefficient between the mean and variance of cells of interest to approximate the Fano factor for that subpopulation. This produced very similar results to the original analysis. Mean-matched Fano factor was calculated by downsampling to match the distributions of firing rates across timepoints, to remove any effect on the Fano factor arising from transient overall changes in firing rate across cells. Again, this produced a similar timecourse to the raw Fano factor. All Fano factor and associated mean firing rate plots in Figure 4.6 are shown for the preferred stimulus condition for a given subpopulation. For the nonselective subpopulation, I used included measurements from both easy upper and easy lower trials.

Population decoding using linear classification was done using the Matlab `fitclinear` function, with L1 regularization and 10-fold crossvalidation. I used only sessions with more than 3 units were recorded simultaneously. I included all units in this analysis, including those with refractory period violations >1%. I used single trial firing rates in the epoch of interest (stimulus, movement, or outcome) to classify the current stimulus, future choice, current movement direction, or reward state. I equalized the number of trials across the two conditions I wanted to classify, by randomly subsampling from the full set of trials. I trained and tested the model over 100 repeats of this to ensure full coverage of the trial conditions. For each repeat, I trained the model on 90% of the data, of which 10% was held out for cross validation, and then tested the accuracy of each repeat of the model on the held out 10% of test data. The accuracy on the test set reported here is the mean accuracy across all repeats. The best cell on a given session was defined as the cell with the largest absolute auROC value on the task parameter of interest. This single trial firing rates of this cell were then classified using the same function as for the population decoding.

## 4.8 Discussion

Here I report single unit data primarily from V1, recorded in a visual discrimination task and a non-visually guided viewing task. I used the stimulus-independent viewing task and behavioral insights to compare representations in cells preferring a stimulus feature that is informative for future choice, versus that is not used to determine future choice. I found differences between the stimulus-selective subpopulations in the trial-to-trial variability of their firing rates, measured by the Fano factor, but not in the non-sensory representations present across these populations. Furthermore, these non-sensory representations are distributed and have low correlation with either stimulus selectivity or the behavioral relevance of the preferred stimulus.

The lack of specific structure for representation of nonsensory variables relative to the known organization of V1 could slightly bolster the hypothesis that the arrival of these signals in V1 serves no specific function, and may trivially reflect broad distribution of generally salient signals across the brain, by virtue of complex long-range connectivity patterns. However, I saw here that movement responsive cells, for example, do not faithfully represent the pure motor content of an action (e.g. direction of movement) but are instead modulated by task epoch, or perhaps the purpose of the present movement (to collect reward versus to initiate a goal-directed sequence of actions by initiating a trial). Whether this modulation occurs locally in V1 circuits or these signals arrive in V1 already with this information content, it suggests that simple corollary discharge from motor effector cells are unlikely to explain the full diversity of the motor responses in V1. In addition, I found that for a substantial number of cells, direction-selective responses begin to emerge while the animal is still stationary; this further suggests that sensory reafference cannot be driving all of the movement-related responses in V1. Finally, the maintenance of information about the previous choice within the population well into the next trial further suggests that movement-related activity in V1 is unlikely to be merely a copy of current motor commands or the resulting sensory changes. It will be interesting to consider whether other frameworks of interpreting firing patterns in V1, such as those that consider V1 activity as generating and updating

a sensory prediction, would be able to better understand the diversity of movement-related activity observed here. This will require more careful experimental design to isolate and probe differential sensory predictions at different timepoints in a trial.

I found that the trial-to-trial variability evolved across sessions differently depending on the stimulus tuning of cells in V1 only during performance of the discrimination task. Specifically, Fano factor was reduced upon stimulus onset to a greater extent for lower-preferring cells, relative to both upper-preferring and untuned cells, and gradually increased in upper-preferring cells over the stimulus period. The selective sustained decrease in variability in lower-preferring cells was less pronounced in preliminary data acquired at a more anterior recording site that may be dominated by units sampled from "higher" cortical regions such as V2 or PPC. In a "passive" stimulus viewing variant of the task, the same relationship was not seen between lower- and upper-preferring cells. In fact, as the duration over which variance was calculated increased, all cells showed a sustained decrease in variability relative to the pre-stimulus period. This was consistent with previously reports, both from the attention literature (Cohen and Maunsell, 2009) and elsewhere (Churchland et al., 2010), which found that stimulus delivery triggers robust quenching of variability across all cells, even in animals that are disengaged or anaesthetized. It is therefore intriguing as to why, in a task that requires stimulus information to guide future decisions, this effect is no longer preserved for the upper-preferring cells, which have been shown in the previous chapters to be more or less ignored during goal-directed behavior.

Based on estimates of the trajectory of the recording electrodes over days, I expect this dataset to be dominated by neurons residing in layers 4, 5, and 6. Therefore, how these data relate to activity recorded via two-photon calcium imaging is an open question, considering these studies primarily sample layer 2/3 neurons, and further, to what extent are differences due to calcium imaging vs electrical recordings, and what is due to differences in layer-to-layer processing. For example, it is known that layer 5 neurons in V1 have larger and more complex-like receptive fields (e.g. wider orientation tuning curves, Niell and Stryker, 2008), and it has also been hypothesized that layer 5 neurons may carry an internal representation of sensory expectations, which may then be compared with feedforward input to drive activity

in layer 2/3 "mismatch" detection neurons (Keller and Mrsic-Flogel, 2018). The asymmetric sampling across cortical layers may therefore skew the proportions of selective cells reported here, but also adds valuable insight into the contributions of deep cortical layer single neurons to a field that has largely relied on imaging techniques restricted to layer 2/3.

# Chapter 5

# Probing the causal role of V1 in the cloud of dots task

## 5.1 Introduction

I sought to probe the functional role of the complex overlapping task representations in V1 by directly manipulating V1 activity during the cloud of dots task. Primarily, I was interested in the causal role of the stimulus representations in V1 in this task. It was previously shown in an auditory analog of this task that stimulation of the direct projection from primary auditory cortex to striatum biased choices towards the tuning preference of the local stimulation site (Znamenskiy and Zador, 2013). Due to the asymmetric and more distributed projection from V1 to striatum that was reported in Chapter 3, I chose not to specifically stimulate corticostriatal projection neurons. Instead, I chose to remain agnostic to the output pathways from V1 that carry out behavioral choice, and carry out bilateral inactivations within V1 itself. However, I kept inactivations localized in order to test the hypothesis that inactivations would bidirectionally influence behavior by "subtracting" information from either the upper visual field or the lower visual field.

The design of this experiment is inspired by the microstimulation experiments done in primates performing the random dot motion (RDM) task. Microstimulation in sensory area MT, which represents motion direction, biases reports toward the preferred direction of the stimulated cluster of cells (Salzman

et al., 1992). The local firing rate increase in these directionally tuned cells is thought to increase evidence sent to downstream circuits about the presence of motion in their preferred direction. However, this experiment required moving the stimulus to the location of the stimulated cells' receptive field, and tuning the target direction of the stimulus to the preference of the neuron. Here, I exploit the retinotopic organization of V1 to inhibit, rather than stimulate, activity in populations of cells encoding stimuli in the upper versus lower field. Here, I use cre-inducible channelrhodopsin-2 (ChR2) specifically expressed in parvalbumin (PV) interneurons to disrupt local activity. Inactivation of activity within cortical areas through activation of GABAergic inhibitory neurons has been successful in both sensory and nonsensory cortices (Rossi et al., 2012, Liu et al., 2014, Zatka-Haas et al., 2019). By inhibiting activity in local subpopulations in binocular V1, I expect to decrease the amount of evidence for stimuli in their preferred region of visual space, and thereby bias behavior towards reports that there are more dots in the opposing region of visual space.

To this end, I targeted different subregions of V1 along the anteroposterior (AP) axis, which has been previously mapped to the elevational retinotopic axis in rats (Gias et al., 2005) and mice (Schuett, Bonhoeffer, and Hübener, 2002). Briefly, I delivered cre-dependent ChR2 bilaterally via viral injection in PV-cre transgenic rats, and directly implanted optical fibers above each injection site (Figure 5.1A). I then stimulated cells expressing ChR2 using a 473nm laser during either the stimulus or outcome epoch of the task. I carried out these manipulations in the version of the task described in Chapter 4.1, where both upper and lower visual streams are integrated (albeit unequally) into the decision. Preliminary results show that local V1 activity during the stimulus period is involved in producing the eventual choice, while little effect is seen following silencing of V1 activity during the outcome epoch.

## 5.2 Inactivation of V1 during stimulus epoch produces choice bias

I first excited ChR2 expressed in PV+ cells to suppress local V1 activity during the stimulus presentation epoch (Figure 5.1B). In the following experiments, I chose to deliver a single sustained pulse for the duration of the stimulus epoch following a short parameter search while recording nearby cells with

tetrodes. The response of an example cell in the vicinity of a sustained 500ms laser pulse is shown in Figure 5.1B. Additionally, I validated that viral delivery was restricted to PV cells by immunostaining for PV (Figure 5.1C).

When I bilaterally stimulated PV cells centered at bregma -5.6mm, suppressing local activity in a retinotopic site preferring lower visual field stimuli, I observed a shift in behavior, with the animal making fewer reports of seeing more dots in the lower hemifield at all stimulus strengths (Figure 5.1D). Meanwhile, behavior on control trials during stimulation sessions was consistent with behavior on the preceding no-stimulation sessions (Figure 5.1D inset), suggesting the effect of the laser was isolated to laser-on trials, and did not affect behavior on the remainder of the session. Conversely, for a stimulation site centered at a more posterior location of bregma -6.0mm, putatively suppressing more upper visual field preferring cells, I observed a behavioral shift in the opposite direction, such that the animal was more likely to report seeing more stimuli in the lower hemifield for all stimulus strengths. These results are consistent with reducing the amount of information carried in V1 about the inactivated subregion, causing the animal to report that the stimulus is dominated by the unmanipulated opposing subregion.

Single session results were largely consistent with the pooled behavior traces shown in Figure 5.1D-E. The effect size between laser-on and interspersed control trials, quantified as the change in the bias term between the two psychometric curves, was significantly different between upper subregion targeted inactivation and lower subregion targeted inactivation subjects (unpaired t-test, p= 0.000167, Figure 5.1F). The comparison between the lower subregion inactivation and the eYFP control group was just significant at p=0.0499, while the comparison between the upper subregion inactivation and the eYFP control group did not reach significance (p=0.862).

Finally, though I targeted two discrete sites for inactivation, I wondered if true injection site position could better explain effect sizes across animals. I found that for both stimulation animals and control animals, there was a correlation between true injection site position and the session-by-session effect size (Figure 5.1G). The correlation between inactivation site location and effect size was significant for ChR2-injected animals (r=-0.7029, p=0.000765), but not significant for control animals (r=-0.7710,

p=0.1193). However, the presence of the same stimulation site-dependent trend for control animals may point to local nonspecific effects of the laser (e.g. heating), rather than visual detection of the laser onset.

## 5.3 Preliminary data suggest outcome epoch inactivations produce no effect on next trial choice

Preliminary data from one animal in which laser stimulation was delivered in the outcome epoch suggests that, in contrast to stimulus epoch inactivation, the effects of inactivation during reward or punishment delivery have minimal effects on next trial behavior. Here, I delivered bilateral stimulation to PV neurons from the time of outcome (e.g. reward, missed reward, or punishment) onset until the following trial initiation, up to a maximum of 5s, which was equal to the punishment timeout duration (Figure 5.2A). Despite the long sustained inactivation in this epoch, I observed little effect on the following trial's choice behavior (Figure 5.2B) compared to control trials following no stimulation. Further separation of trials by the preceding outcome – that is, the outcome during which laser stimulation was given – yielded little difference between trials following stimulation and control trials. However, comparing within trials following laser stimulation, there is a slight effect of previous outcome on future choice: trials following reward are less accurate than trials following no-reward (Figure 5.2C). This history-dependent effect is not present for control trials. Further investigation is required to test the robustness of this preliminary observation.

## 5.4 Methods

Rats were bred by crossing PV-Cre+ BAC transgenic rats (Strain: LE-Tg (Pvalb-iCre)2Ottc, NIDA/NIMH) with one another or PV-Cre- Long Evans rats. Rats were genotyped upon weaning and PV-Cre+ rats selected for training. Following training, I delivered Cre-dependent Channelrhodopsin-2 bilaterally via

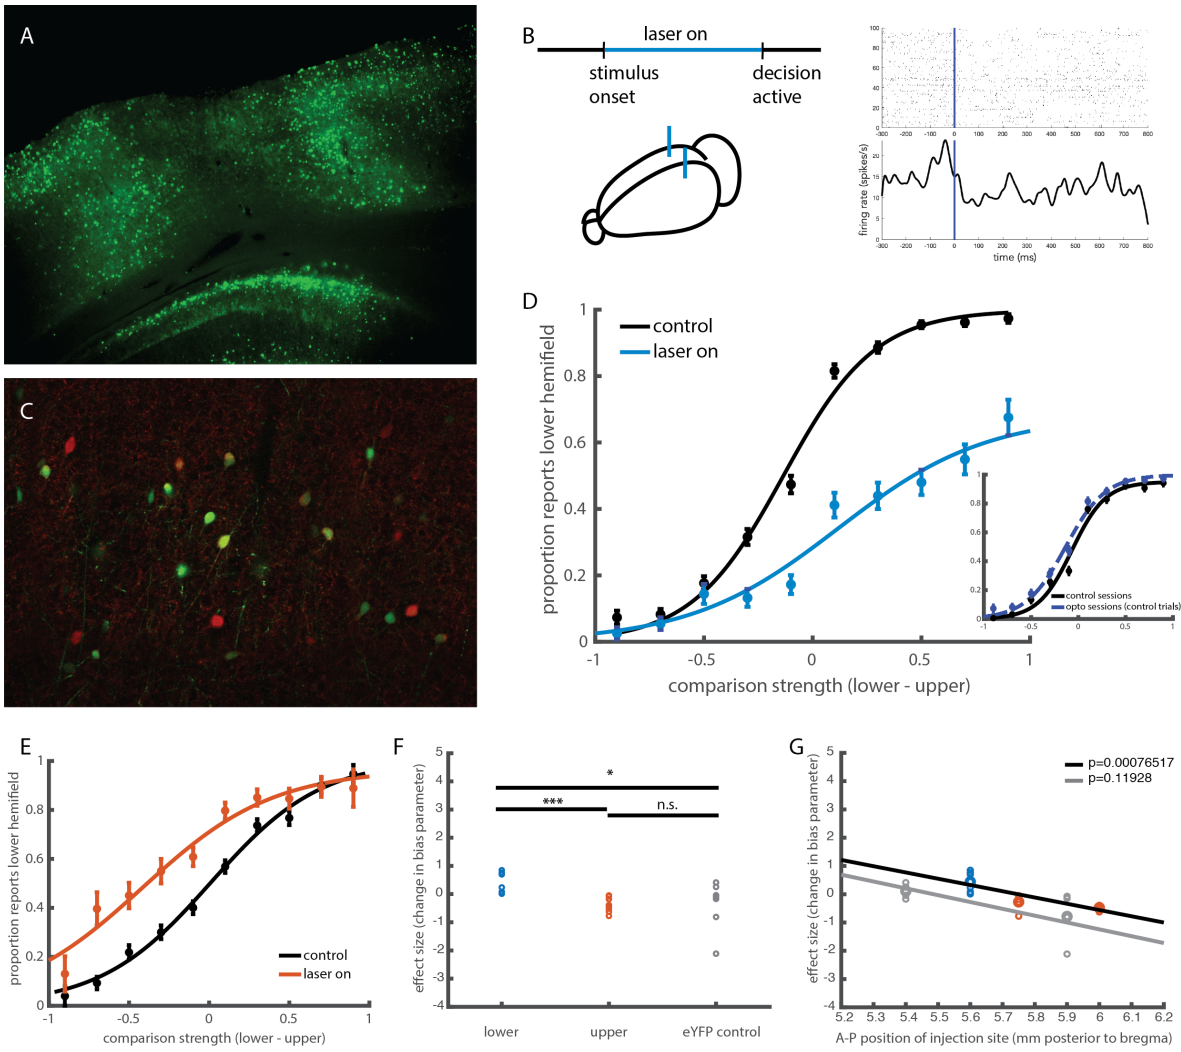
A
B
laser on
stimulus onset
decision active
firing rate (spikes/s)
time (ms)
C
D
control
laser on
proportion reports lower hemifield
comparison strength (lower - upper)
control sessions
opto sessions (control trials)
E
proportion reports lower hemifield
control
laser on
comparison strength (lower - upper)
F
effect size (change in bias parameter)
*
***
n.s.
lower
upper
eYFP control
G
p=0.00076517
p=0.11928
effect size (change in bias parameter)
A-P position of injection site (mm posterior to bregma)

viral vector, and implanted optic fibers directly above the injection site. I targeted bilateral upper ('posterior' V1, bregma -6.1mm AP, +/-4.5ML) and lower (frontal V1, bregma -5.5mm AP, +/-4.5ML) preferring subregions of V1, with coordinates derived from my recording site coordinates and the reference brain atlas. At each site I pressure injected 90nL of either AAV9.hSyn.DIO.ChR2.eYFP or a control virus AAV9.hSyn.DIO.eYFP at 500 μm and 800 μm below the surface of the brain. 250 μm diameter fibers were affixed to the surface of the skull above each injection site, with the tip of the fiber shaved flat and lowered to the surface of the brain.

To excite channelrhodopsin in PV cells, I used two 473nm lasers paired to the same input trigger to deliver light to both implanted fibers simultaneously through coupling through a LC-FC fiber optic cable. I calibrated laser power to deliver 8mW at the tip of each fiber optic. Given light loss at the interface with the implanted fibers, I expected approximately 5-6mW laser power to be delivered at the brain surface. I used PulsePal (Sanders and Kepecs, 2014) in pulse-gated mode to communicate between the behavior state machine and the laser. This allowed me to turn on the laser during the relevant behavioral epochs, and turn it off with fast latency upon stimulus abortion or next trial initiation, for example.

During a given stimulation session, 30% of trials were paired with laser, while the remaining trials

FIGURE 5.1 *(preceding page)*: **PV-mediated inactivation of V1 during stimulus epoch affects choice behavior:** A. Expression of eYFP-tagged ChR2 at injection site in frontal V1 (bregma -5.6mm). B. Left: Schematic showing bilateral stimulus period inactivation. A sustained laser pulse was delivered from stimulus onset until decision tone onset ( 650ms). Right: Raster and PSTH of example cell recorded during 500ms sustained laser pulse delivery, aligned to laser onset. C. PV immunostaining (Alexa-647, in red) colocalizes with cre-dependent eYFP labeled cells. D. Pooled psychometrics for one animal with frontal V1 inactivation for laser-on versus control trials (n=6 sessions). Inset panel shows psychometric behavior on control trials during stimulation sessions is consistent with baseline behavior measured on the preceding no-stimulation session. E. Same as D, but for an animal with a more posterior V1 injection and implant site. F. Summary of effect size (measured as bias parameter difference between laser-on and control trials on individual sessions) for each of 3 groups of animals: lower subregion (frontal V1, N=1), upper subregion (more posterior V1, N=2), and an eYFP positive control group (N=2). The effect size between upper and lower inactivation groups is significantly different (unpaired t-test, p = 0.000167). The effect size between lower inactivation and eYFP. control groups is also significant (p=0.0499). G. Summary of effect size, measured as in F, as a function of the injection site position. Larger points show mean effect size over sessions. ChR2-injected subjects are shown in blue and orange, color conventions as in F. Control eYFP-injected animals are shown in grey. Linear regressions show a correlation between AP-position of injection site and effect size in ChR2-injected subjects (r=-0.7029) and though not significant, a similar trend in control subjects (r=-0.7710).

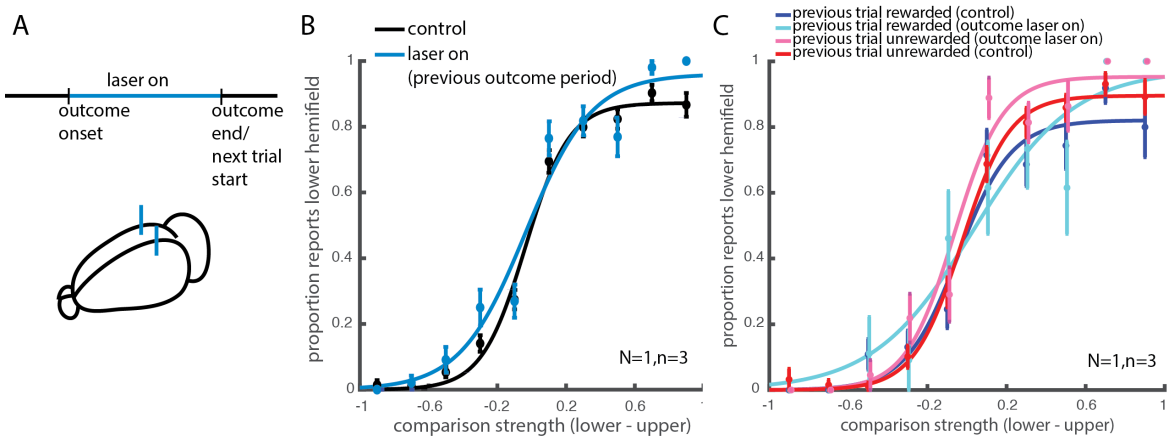


FIGURE 5.2: **Inactivation of V1 during outcome epoch has little effect on next trial choice:** A. Schematic showing bilateral outcome period inactivation. A sustained laser pulse was delivered from outcome onset until either the end of the punishment duration (5s) or initiation of the next trial, whichever came first. B. Pooled psychometrics for one animal with frontal V1 inactivation, comparing choice on trials that were preceded by outcome-epoch inactivation to control trials (n=3 sessions). C. Psychometrics in B, further split by outcome of the previous trial.

were used as control trials. Only one trial a week included laser stimulation; however, animals were tethered with the fiber optic cable on all sessions. To interfere with the animal's ability to detect when the laser was on, all sessions also included an overhead masking light.

## 5.5 Discussion

Other optogenetics studies have identified a role for normal V1 activity in detecting contralateral stimuli (Burgess et al., 2017), detecting differences and changes in orientation and contrast (Glickfeld, Histed, and Maunsell, 2013, Resulaj et al., 2018), and instructing visually-guided action timing (Namboodiri et al., 2015). Here I tested the effect of disrupting activity in tuned subregions on the associated choice, both within and across trials. While more subjects and sessions will be necessary to make conclusions on the nature and strength of the correlations between inactivation site and behavioral effect, these preliminary results begin to add to our view that despite complex activity dynamics over the course of the trial, it is stimulus epoch V1 activity that is directly causal for visually-based decisions. Furthermore, these results

are consistent with a model in which inhibition of activity in a spatially tuned subpopulation decreases the amount of evidence for that region of visual space, and thereby increases the proportion of choices for the opposing subregion. While stimulus period inactivation had a bidirectional effect on upcoming choice, depending on the local tuning of the inactivation site, preliminary results showed outcome period inactivation had no such effect on future choice.

Here, in order to probe a bidirectional shift in the choice behavior, I carried out these manipulations in the context of a task where upper and lower streams of information were independently and uniformly drawn. I later began to understand that under these stimulus conditions, animals continued to overweight the lower information stream, sometimes to varying degrees over sessions, potentially adding ambiguity to the interpretation of pooled results. Estimates of the degree to which animals overweight the lower stream can be noisy in single sessions, and therefore studying the effects of variable strategy requires more data points than is presented here. Moving forward, I will investigate whether effect size correlates with the session-by-session strategy of the animal. In addition, I will repeat these experiments in animals doing a true comparison between the upper and lower visual streams of information. It will be interesting to see if the preliminary observation in Figure 5.1F-G of smaller absolute effect sizes for upper compared to lower visual field inactivations is due to the imbalance in stimulus usage by these animals, or simply sampling noise.

Interpretation of a bidirectional shift, especially in the case of inactivations that selectively increase the probability of choosing the opposing choice, suggests a race or selection decision model that arbitrates between two streams of inputs. This has been explored extensively in the macaque RDM literature, where reaction time tasks provide an extra parameter against which to test the effects of microstimulation on the decision process (Ditterich, Mazurek, and Shadlen, 2003). This fixed time task does not allow us to judge temporal signatures of the decision process related to inactivations.

The preliminary lack of an effect of outcome epoch inactivation requires further exploration. If this (lack of) effect persists with more repeats of this experiment, there are several possible explanations. The first is that outcome period activity (which dominates the firing of many single cells across the

population, as shown in the previous chapter) reflects information that is not integrated locally into V1 circuits, or read out from downstream circuity involved in visually-guided choice. The second reason is that while these inactivations were spatially optimized to differentiate between populations preferring different regions of visual space, they do not lie along the correct axis for preferentially targeting outcome activity, and therefore any effect of the PV-inactivation may be cancelled out across the local population. Indeed, I was not able to find a clear relationship between outcome selectivity and spatial tuning in V1 neurons (Chapter 4) in the original version of this task. The third possible reason is that the efficacy of the PV stimulation has a proportionally smaller effect on outcome period activity, either due to sustained inactivation leading to adaptation, or simply being too low in magnitude or sparsely delivered to PV cells to effectively silence outcome responses or periods of higher activity. In contrast, a study inactivating mouse V1 found that altered visual representations influence behavioral output even with only moderate changes in neuronal activity (Glickfeld, Histed, and Maunsell, 2013). These open questions can be tackled by conducting further tetrode recordings during optogenetic stimulation.

# Chapter 6

# Discussion

## 6.1 Summary

Here, I have presented behavioral, neural, and causal data on how visual information available to an animal is used to drive behavior. First, I developed a novel visually-guided discrimination task along with a software implementation of head position control based on online video tracking and closed loop behavioral control. Then, I showed that while animals can learn to compare the stimulus strength between two subregions of the stimulus (upper versus lower) when the specific statistics of the stimulus environment over time dictate it, they typically converge on using solely the information in the lower visual field, both when this is equally optimal to the comparison strategy, and when it is slightly suboptimal.

I showed that single neurons in primary visual cortex are selective to regions in visual space, with some neurons firing preferentially when stimuli predominantly appeared in the upper half, while others fired preferentially when more stimuli appeared in the lower half. In this dataset, these tuned subpopulations were roughly equal in size. I showed that the stimulus preference of a neuron had at most weak correlations with its choice probability, and with its representations of other task features in other epochs of the task, such as choice direction and reward. Indeed, individual neurons sometimes preferred the opposing choice to the one associated with their preferred stimulus. While tuning to individual stimulus features was poorly correlated with the behavioral use of the preferred stimulus, the statistics of stimulus evoked firing showed significant differences both within and between tasks. Specifically, the variability

of firing over trials showed a sustained decrease for the subpopulation preferring the "useful" lower portion of the stimulus, while non-selective and upper-preferring cells were more variable in their firing on average, and gradually increased this variability over the stimulus period. In contrast to this, the firing rate variability was more similar across all populations of cells in a task variant where all visual information was by definition "useless" in guiding choice behavior. Here, all tuning-defined subpopulations showed similar sustained decreases in variability for the duration of the stimulus.

Finally, I showed that local stimulus representations in primary visual cortex may have a causal effect on choice behavior, such that selectively inactivating a portion of the retinotopic map in V1 during stimulus presentation reduces behavioral reports that there are more stimuli in that region. Meanwhile, preliminary results show little effect of inactivation of outcome epoch activity on the next trial behavior of the animal.

## 6.2 Perspectives on behavior

The behavioral results point to a consistent asymmetry in stimulus use across visual space that has not been reported in the context of a discrimination task before, though innate fear responses have suggested there are intrinsic differences in how a stimulus (e.g. an expanding circle, or "looming" stimulus) may be interpreted depending on where it is delivered (from above or from the side, Yilmaz and Meister, 2013, Wallace et al., 2013). This asymmetry looks at the behavioral level like an extreme form of attention, whereby one stream of information (the "attended" stream) strongly or exclusively determines the behavior of the animal, regardless of what appears in the opposing "distractor" stream. Here, however, these designations are not flexibly assigned by the experimenter on a trial-by-trial basis, but instead appears intrinsic and consistent across all animals. To what extent does this bias reflects a true attentional process versus an innate processing pattern, and how are similarities and differences between these processes reflected in the computations that arise in sensory cortex? Attentional cues typically act over short time scales, whereas the behavioral asymmetry observed here only varied over many hundreds of trials in a history-dependent manner. It is interesting that we observed a similar variability signature in the

primary visual cortex of these animals to what has been reported in (typically higher) visual cortices during attention; further work comparing additional features of neural activity may provide more insight into the relationship between these behavioral processes.

The mechanisms underlying this intrinsic and persistent bias towards the lower visual field are unknown and intriguing. Differences in stimulus use may be mediated by any number of long-range projection, top-down modulation, or local processing differences, one candidate for which was presented here in the form of an asymmetric projection to dorsomedial striatum. It has been proposed that anatomical connectivity, set up through evolution, may serve to expedite or constrain animals' learning or performance during its lifetime (Zador, 2019). It would be interesting to understand the effect of manipulations of the projection asymmetry (e.g. through pathway-specific ablation) on strategy choice and learning dynamics on this task.

Beyond the specific implementation of this behavioral asymmetry, these behavioral data also suggest two important things about strategy learning. The first is that categorical shifts in strategy (i.e. changes between ignoring one stream of information versus integrating it into a comparison metric) can be learned from long timescale statistics of the sensory environment. Here, a change in the stimulus distribution over trials, but not the structure of stimuli on individual trials, was sufficient to cause the animals to adjust their behavioral strategy from the suboptimal strategy, to the optimal one. The second feature of strategy learning that can be inferred from this data is that strategy exploration in a uniform landscape is not the same as strategy learning due to external pressures. While animals transiently used at least a partial comparison strategy in early learning, this was not sustained for longer than 1 week for any animal tested here. However, once animals were compelled to adopt a comparison strategy due to the stimulus distribution, this strategy consistently remained stable for at least 3 weeks after stimulus distribution was returned to a uniform distribution. This may be due to increased training days during which the comparison strategy remained much more optimal than the alternative strategy, and therefore providing continuous reinforcement of this strategy. Related to this, I observed in some animals a slight overfitting of the behavior such that they performed better on "adversarial" trials than on baseline trials, though this

was quick to drop off as soon as stimulus distribution went back to uniform.

One would like to understand the source of the intrinsic, seemingly suboptimal bias for the lower-visual field, and why this occurs despite implementation of the "correct" strategy during early learning on the uniform distribution version of the task. There are possible ethological (motivational salience of the upper versus lower visual field, with threatening stimuli likely to occur in the upper field, and stimuli useful for exploratory or foraging behaviors appearing predominantly in the lower visual field) or physical (with real-world illumination coming from the sun, proportions of luminance-on cells or thresholds might differ across the elevational extent of the rodent visual field) reasons for this bias, with potential implementation through asymmetric circuitry or local differences in processing. However, if we consider this task as a goal-directed behavior, the apparent preference for suboptimality is puzzling. For an ideal observer facing a uniform distribution of stimuli drawn independently across the upper and lower visual streams, the lower-thresholding strategy leads to a 25% loss in reward collection relative to the comparison strategy. Of course, animals are not ideal observers. Noise at both the sensory periphery and downstream areas limits the accuracy of the animal. Depending on the mechanism by which comparison is carried out, integrating information across both stimulus streams could lead to additive noise, thereby decreasing the accuracy of the animal across trials. This suggests that in a stimulus design where the information in the upper and lower streams are perfectly correlated, the lower-thresholding strategy is at worst equally optimal to the comparison strategy, and at best, superoptimal compared to the comparison strategy.

Another mechanism by which noise can decrease the accuracy of the animal is through scalar variability – the scaling of noise with the number of stimulus pulses – which has been found to flatten the psychometric curve for trials with large numbers of total stimuli in an auditory click counting task (Scott et al., 2015). Here, the behavior in the cloud of dots task with independently drawn stimulus streams did not show a dependence of behavioral sensitivity on total stimulus number, especially once animals converged on a lower-thresholding strategy. However, when animals were trained to stably implement a comparison strategy, a small difference can be observed in the sensitivity of the behavior between trials

with fewer and trials with more dots (Figure 3.2). This may be another source of noise that constrains the optimality of the comparison strategy. Further behavioral experiments involving changes to the stimulus distribution and maximum reward rate at different points in training may help to identify whether animals indeed have and/or are influenced by subtle sensitivities to different sources of noise.

Many questions remain about the strategy learning observed here, including whether animals store a stimulus history or inferred generative stimulus distribution which they then use to guide this learning by comparing with past choices and outcomes, or whether they reduce these changes down to a simpler cached value, such as reward rate per unit time. Another open question is when animals are in the intermediate stages between the two extreme strategies I measure here, does their behavior consist of switching between strategies across trials, or is the appearance of an intermediate strategy due to intermediate weighting of the two stimulus streams within trials, or another strategy altogether? Of course, the neural circuit underpinnings of these adaptive changes are unknown, as are the ways in which local computations in V1 may reflect strategy changes.

## 6.3 Perspectives on neural representations

To probe the neural underpinnings of this behavior, I recorded neural activity in V1 with the goal of comparing information content and firing statistics within and between tasks. I first sought to characterize the diversity of responses in V1 single neurons during a visually-guided behavior. The design of this task, with clearly delineated task epochs, allowed us to study the features of single neuron responses during the sensory, motor, and outcome epochs individually, and map how they relate to each other across the population. Similar to previous reports, I found prominent movement epoch responses across the population. Contrary to some reports (Steinmetz et al., 2019), I found that a majority of cells had a direction preference in their movement epoch activity, with 14% of these cells showing direction selectivity prior to movement onset. In addition, I found that there was variability in single neuron encoding of the movement direction dependent on the task epoch or precise ports traversed. These two variables are of course correlated, as all choice movements originate from the center port, while all initiation movements

originate from the side ports. The result is a specific encoding of movement parameters by these neurons beyond initiation or even direction. Possible features represented in these responses might include spatial location, abstract goal, or value expectation. These observations together suggest that movement responses in primary visual cortex cannot be due purely to sensory reafference or corollary discharge of motor command neurons, but that they instead either receive or integrate locally other task features. In general, the diversity of individual and combinations of features shown to be encoded in V1 neurons here bears some resemblance to the kind of mixed or distributed selectivity that has been reported in higher cortices such as prefrontal and parietal cortices (Rigotti et al., 2013, Raposo, Kaufman, and Churchland, 2014).

I interrogated the hypothesis that the behavioral relevance of stimulus-tuned neurons might influence the amount of information encoded by the same neurons about other features of the task. I reasoned that cells responding more when the stimulus is weighted strongly towards the lower half of the stimulus would be more relevant to a choice that is based on the strength of the stimulus in that half. One might therefore expect that such cells would have higher choice probabilities, or carry more information about choice direction during the movement epoch, or outcome during the outcome epoch. Instead, I found that there was little difference in the distribution of choice direction information represented in cells that preferred upper (unused) versus lower (used) stimuli. In addition, while there were significant differences between the two stimulus-tuned subpopulations in their choice probability and reward selectivity distributions, a look at the spread of these distributions showed that in neither of these cases did the lower-preferring subpopulation have a higher proportion of highly selective cells for choice probability or reward delivery. Instead, differences in the distribution were more likely due to offset of the distribution peaks between the two subpopulations.

These data, along with the weak correlations across the population between the encoding of these features within single neurons, suggests that encoding of nonsensory variables in V1 is organized in an orthogonal or random fashion with response to sensory encoding, and is not correlated with goal-directed use of the sensory information represented by those neurons. This contrasted with a previous

report (Poort et al., 2015) that choice signals developed over learning preferentially within neurons that preferred the visual stimulus that led to reward. In this study, animals were asked to discriminate the orientation of a full-field stimulus as they traversed a virtual corridor, and decide whether to lick or withhold licking. The specifics of the task design may account for some of the differences in neural activity seen between this study and the data presented here, such as the use of one target stimulus and one "distractor" or non-target stimulus, or the choice report taking the form of a go-no go. The comparison between a discrimination task based on orientation versus one based on visual space is further complicated by the fact that representations of orientation and visual space are organized differently across V1, and so may also map differently onto other features arising in V1 from top-down projections. One might also speculate that there is a limit, either in terms of behavioral demands or feature dimensionality, beyond which targeted selectivity ("labelled lines") may no longer be computationally efficient (Rigotti et al., 2013).

Meanwhile, I observed modulation of the trial-by-trial variability of responses that differed across subpopulations. Though the effect of a lower Fano factor in the lower-preferring subpopulation matched my expectations based on the attention literature, whereby local spatial attention decreases the variability of firing over trials, I did not expect that this effect would be driven by reduced magnitude and duration of stimulus-evoked quenching in the upper-preferring and nonselective subpopulations. Instead of suggesting the influence of top-down modulation preferentially targeting lower-preferring cells, this raises the possibility that local computations in V1 may change during goal-directed behavior. For example, though upper-preferring cells showed a sustained decrease in Fano factor for the duration of the stimulus period in a stimulus-independent task, their Fano factor instead gradually slightly increased over the stimulus period in this discrimination task. Previously, increases in Fano factor have been found to accompany accumulation processes, as independent noise adds up over the accumulation of momentary evidence (Churchland et al., 2011).

## 6.4 Future directions

Moving forward, these data provide many jumping off points for future inquiry. The asymmetric projection pattern from V1 subregions to dorsomedial striatum mentioned briefly above might be considered in terms of models of corticostriatal pathway function, especially compared to the corticocollicular pathway. Whereas the lower-preferring subregion of V1 dominated the projections to striatum, superior colliculus received approximately equal density of projections from upper- and lower-preferring V1 subregions. Distinct computational roles have been proposed for these two circuits in the context of a drift diffusion model (Lo and Wang, 2006, Crapse, Lau, and Basso, 2018). Furthermore, visual representations in the superior colliculus have been implicated in visual fear responses in rodents (Shang et al., 2018) and attentional modulation in primates (Krauzlis, Lovejoy, and Zénon, 2013). Whether either or both of the differing projection patterns from V1 to striatum and superior colliculus are meaningful for the execution – and manner of execution – of this task can first be tested by selective ablation of one of these pathways, either transiently or chronically. Further experiments might then explore the relative effects of selectively ablating the projection from differently tuned subregions of V1 on task execution or strategy.

A second point that could be well-served by both theoretical and experimental followup is the observation that nonsensory features appear to be organized in a distributed fashion, relative to sensory representations in V1. One possible reason for this is that such an organization might allow flexibility, by allowing cells of all stimulus preferences access to information about relevant behavioral events such as choice and outcome, which could be valuable in order to allow exploration in the context of a changing stimulus environment, for example. This could also help to explain why inactivation of V1 subregions during the outcome epoch, when many neurons fire strongly, has little effect on the next trial in a well-trained animal. Another possibility is that the techniques used here mask a true structure to the organization of nonsensory representations. For example, if nonsensory representations were mapped across visual cortex in either an orthogonal dimension to the elevational retinotopic axis, or if they mapped at a finer repeating scale, similar to orientation pinwheels in the primate and cat, a broad

inactivation along the elevational retinotopic axis would obscure any effects. Similarly, blind sampling of single cells using chronic tetrode implants does not allow comparison of the spatial distribution of these variables relative to one another across V1. Two-photon imaging could be used to gain access to this information, but would be limited to superficial cortical layers, especially in the rat, which may skew the properties of recorded cells.

In conclusion, while rodent visual decision-making is becoming an increasingly powerful paradigm for studying both visual processing and circuits underlying sensorimotor associations, I hope I have shown here that experimenter-designed behavioral rules do not always directly translate to an animal's implementation, which can often be better understood through direct behavioral experiments. I have also shown that the organization of information in visual cortical neurons is not well-correlated between task-defined features in our task, nor is it strongly dependent on the causal behavioral influence of those neurons. However, the information that is present in visual cortex is often highly specific, even at the level of single neurons. What might be the reason for these highly selective yet distributed signals? There are theories that suggest distributed coding of information across areas or even the full brain may help to sustain a global network representation that permits a wide diversity of readout functions (Rigotti et al., 2013, Mashour et al., 2020). Our result here contrasts with several expectation-based behavioral paradigms where representations of correlated task parameters have been reported within the same neurons. The difference in the organization of information across tasks and types of expectations in V1 may serve as a valuable window on the limits of organized representations within a relatively well understood cortical region.

# Bibliography

Akrami, Athena, Charles D. Kopec, Mathew E. Diamond, and Carlos D. Brody (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* **554**.

Armstrong, Katherine M. and Tirin Moore (2007). Rapid enhancement of visual cortical response discriminability by microstimulation of the frontal eye field. *Proceedings of the National Academy of Sciences of the United States of America* **104**: 9499–9504.

Attinger, Alexander, Bo Wang, and Georg B Keller Correspondence (2017). Visuomotor Coupling Shapes the Functional Development of Mouse Visual Cortex. *Cell* **169**: 1291–1302.e14.

Bengio, Yoshua (2017). The Consciousness Prior. *arXiv*.

Brainard, David H. (1997). The Psychophysics Toolbox. *Spatial Vision* **10**: 433–436.

Briggs, Farran, George R Mangun, and W Martin Usrey (2013). Attention enhances synaptic efficacy and the signal-to-noise ratio in neural circuits. *Nature* **499**: 476–80.

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

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

Brody, Carlos D. and Timothy D. Hanks (2016). Neural underpinnings of the evidence accumulator. *Current Opinion in Neurobiology* **37**: 149–157.

Brunton, Bingni W., Matthew M. Botvinick, and Carlos D. Brody (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**: 95–98.

Burgess, Christopher P., Armin Lak, Nicholas A. Steinmetz, Peter Zatka-Haas, Charu Bai Reddy, Elina A.K. Jacobs, Jennifer F. Linden, Joseph J. Paton, Adam Ranson, Sylvia Schröder, Sofia Soares, Miles J. Wells, Lauren E. Wool, Kenneth D. Harris, and Matteo Carandini (2017). High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. *Cell Reports* **20**.

Carandini, Matteo, Jonathan B Demb, Valerio Mante, David J Tolhurst, Yang Dan, Bruno A Olshausen, Jack L Gallant, and Nicole C Rust (2005). Do We Know What the Early Visual System Does? *The Journal of Neuroscience* **25**: 10577–10597.

Celebrini, Simona and William T. Newsome (1994). Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. *Journal of Neuroscience* **14**: 4109–4124.

Churchland, Anne K, R Kiani, R Chaudhuri, Xiao-jing Wang, Alexandre Pouget, and M N Shadlen (2011). Article Variance as a Signature of Neural Computations during Decision Making. *Neuron* **69**: 818–831.

Churchland, Mark M, Byron M Yu, John P Cunningham, Leo P Sugrue, Marlene R Cohen, Greg S Corrado, William T Newsome, Andrew M Clark, Paymon Hosseini, Benjamin B Scott, David C Bradley, Matthew A Smith, Adam Kohn, J Anthony Movshon, Katherine M Armstrong, Tirin Moore, Steve W Chang, Lawrence H Snyder, Stephen G Lisberger, Nicholas J Priebe, Ian M Finn, David Ferster, Stephen I Ryu, Gopal Santhanam, Maneesh Sahani, and Krishna V Shenoy (2010). Stimulus onset quenches neural variability : a widespread cortical phenomenon. *Nature Neuroscience* **13**: 369–378.

Cohen, Marlene R and Adam Kohn (2011). Measuring and interpreting neuronal correlations. *Nature Publishing Group* **14**: 811–819.

Cohen, Marlene R. and John H.R. Maunsell (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience* **12**: 1594–1600.

Cook, Erik P. and John H.R. Maunsell (2002). Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nature Neuroscience* **5**: 985–994.

Cossell, Lee, Maria Florencia Iacaruso, Dylan R. Muir, Rachael Houlton, Elie N. Sader, Ho Ko, Sonja B. Hofer, and Thomas D. Mrsic-Flogel (2015). Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* **518**.

Crapse, Trinity B., Hakwan Lau, and Michele A. Basso (2018). A Role for the Superior Colliculus in Decision Criteria. *Neuron* **97**: 181–194.e6.

Ding, Long and Joshua I. Gold (2010). Caudate encodes multiple computations for perceptual decisions. *Journal of Neuroscience* **30**: 15747–15759.

— (2012). Separate, Causal Roles of the Caudate in Saccadic Choice and Execution in a Perceptual Decision Task. *Neuron* **75**: 865–874.

Ditterich, Jochen, Mark E. Mazurek, and Michael N. Shadlen (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience* **6**: 891–898.

Engel, Tatiana A, Nicholas A Steinmetz, Marc A Gieselmann, Alexander Thiele, Tirin Moore, and Kwabena Boahen (2016). Selective modulation of cortical state during spatial attention. *Science* **354**: 1140–1144.

Felleman, Daniel J. and David C. Van Essen (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* **1**: 1–47.

Fiser, Aris, David Mahringer, Hassana K. Oyibo, Anders V. Petersen, Marcus Leinweber, and Georg B. Keller (2016). Experience-dependent spatial expectations in mouse visual cortex. *Nature Neuroscience* **19**.

Gavornik, Jeffrey P. and Mark F. Bear (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nature Neuroscience* **17**: 732–737.

Gias, C., N. Hewson-Stoate, M. Jones, D. Johnston, J. E. Mayhew, and P. J. Coffey (2005). Retinotopy within rat primary visual cortex using optical imaging. *NeuroImage* **24**: 200–206.

Girman, S. V. (1985). Responses of neurons of primary visual cortex of awake unrestrained rats to visual stimuli. *Neuroscience and Behavioral Physiology* **15**: 379–386.

Glickfeld, Lindsey L., Mark H. Histed, and John H.R. Maunsell (2013). Mouse primary visual cortex is used to detect both orientation and contrast changes. *Journal of Neuroscience* **33**: 19416–19422.

Glickfeld, Lindsey L. and Shawn R. Olsen (2017). Higher-Order Areas of the Mouse Visual Cortex. *Annual Review of Vision Science* **3**: 251–273.

Glickfeld, Lindsey L., Mark L. Andermann, Vincent Bonin, and R. Clay Reid (2013). Cortico-cortical projections in mouse visual cortex are functionally target specific. *Nature Neuroscience* **16**: 219–226.

Gold, Joshua I. and Michael N. Shadlen (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience* **30**: 535–574.

Gottfried, Jay A., John O'Doherty, and Raymond J. Dolan (2003). Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science* **301**: 1104–1107.

Guitchounts, Grigori, Javier Masis, Steffen BE Wolff, and David Cox (2020). Encoding of 3D Head Orienting Movements in Primary Visual Cortex. *bioRxiv*: 2020.01.16.909473.

Han, Yunyun, Justus M. Kebschull, Robert A.A. Campbell, Devon Cowan, Fabia Imhof, Anthony M. Zador, and Thomas D. Mrsic-Flogel (2018). The logic of single-cell projections from visual cortex. *Nature* **556**.

Hanks, Timothy D., Charles D. Kopec, Bingni W. Brunton, Chunyu A. Duan, Jeffrey C. Erlich, and Carlos D. Brody (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**: 220–223.

Haxby, James V, Cheryl L Grady, Barry Horwitz, Leslie G Ungerleider, Mortimer Mishkin, Richard E Carson, Peter Herscovitch, Mark B Schapiro, and Stanley I Rapoport (1991). Dissociation of object and spatial visual processing pathways in human extrastriate cortex (regional cerebral blood flow/positron emission tomography). *Proc. Nati. Acad. Sci. USA* **88**: 1621–1625.

Hembrook-Short, Jacqueline R., Vanessa L. Mock, and Farran Briggs (2017). Attentional Modulation of Neuronal Activity Depends on Neuronal Feature Selectivity. *Current Biology* **27**: 1878–1887.e5.

Herrero, Jose L, Marc A Gieselmann, Mehdi Sanayei, and Alexander Thiele (2013). Attention-Induced Variance and Noise Correlation Reduction in Macaque V1 Is Mediated by NMDA Receptors. *Neuron* **78**: 729–739.

Hubel, David H. and Torsten N. Wiesel (1972). Laminar and columnar distribution of geniculo-cortical fibers in the macaque monkey. *Journal of Comparative Neurology* **146**: 421–450.

Hubel, D.H. and T.N. Wiesel (1959). Receptive Fields of Single Neurones in the Cat's Striate Cortex. *Journal of Physiology* **148**: 574–591.

Hunnicutt, Barbara J., Bart C. Jongbloets, William T. Birdsong, Katrina J. Gertz, Haining Zhong, and Tianyi Mao (2016). A comprehensive excitatory input map of the striatum reveals novel functional organization. *eLife* **5**.

Jaramillo, Santiago and Anthony M. Zador (2011). The auditory cortex mediates the perceptual effects of acoustic temporal expectation. *Nature Neuroscience* **14**: 246–253.

Keller, Georg B., Tobias Bonhoeffer, and Mark Hübener (2012). Sensorimotor Mismatch Signals in Primary Visual Cortex of the Behaving Mouse. *Neuron* **74**: 809–815.

Keller, Georg B. and Thomas D. Mrsic-Flogel (2018). Predictive Processing: A Canonical Cortical Computation. *Neuron* **100**.

Khibnik, Lena A, Nicolas X Tritsch, and Bernardo L Sabatini (2014). A Direct Projection from Mouse Primary Visual Cortex to Dorsomedial Striatum. *PLoS ONE* **9**.

Kim, J N and M N Shadlen (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature neuroscience* **2**: 176–85.

Kleinfeld, David, Ehud Ahissar, and Mathew E. Diamond (2006). Active sensation: insights from the rodent vibrissa sensorimotor system. *Current Opinion in Neurobiology* **16**: 435–444.

Krauzlis, Richard J., Lee P. Lovejoy, and Alexandre Zénon (2013). Superior Colliculus and Visual Spatial Attention. *Annual Review of Neuroscience* **36**: 165–182.

Krumin, Michael, Julie J. Lee, Kenneth D. Harris, and Matteo Carandini (2018). Decision and navigation in mouse parietal cortex. *eLife* **7**.

Leinweber, Marcus, Daniel R. Ward, Jan M. Sobczak, Alexander Attinger, and Georg B. Keller (2017). A Sensorimotor Circuit in Mouse Cortex for Visual Flow Predictions. *Neuron* **95**.

Licata, Angela M., Matthew T. Kaufman, David Raposo, Michael B. Ryan, John P. Sheppard, and Anne K. Churchland (2017). Posterior parietal cortex guides visual decisions in rats. *Journal of Neuroscience* **37**: 4954–4966.

Liu, Ding, Xiaowei Gu, Jia Zhu, Xiaoxing Zhang, Zhe Han, Wenjun Yan, Qi Cheng, Jiang Hao, Hongmei Fan, Ruiqing Hou, Zhaoqin Chen, Yulei Chen, and Chengyu T. Li (2014). Medial prefrontal activity during delay period contributes to learning of a working memory task. *Science* **346**: 458–463.

Lo, Chung Chuan and Xiao Jing Wang (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nature Neuroscience* **9**: 956–963.

Lopes, Gonçalo, Niccolò Bonacchi, João Frazão, Joana P Neto, Bassam V Atallah, Sofia Soares, Luís Moreira, Sara Matias, Pavel M Itskov, Patrícia A Correia, Roberto E Medina, Lorenza Calcaterra, Elena Dreosti, Joseph J Paton, and Adam R Kampff (2015). *Bonsai: an event-based framework for processing and controlling data streams*.

Mashour, George A., Pieter Roelfsema, Jean Pierre Changeux, and Stanislas Dehaene (2020). Conscious Processing and the Global Neuronal Workspace Hypothesis. *Neuron* **105**: 776–798.

Mathis, Alexander, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience* **21**: 1281–1289.

Maunsell, John H.R. (2015). Neuronal Mechanisms of Visual Attention. *Annual Review of Vision Science* **1**: 373–391.

Meyer, Arne F., John O'Keefe, and Jasper Poort (2020). Two distinct types of eye-head coupling in freely moving mice. *bioRxiv*: 2020.02.20.957712.

Mishkin, Mortimer, Leslie G. Ungerleider, and Kathleen A. Macko (1983). Object vision and spatial vision: two cortical pathways. *Trends in Neurosciences* **6**: 414–417.

Mitchell, Jude F, Kristy A Sundberg, and John H Reynolds (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* **55**: 131–41.

— (2009). Article Spatial Attention Decorrelates Intrinsic Activity Fluctuations in Macaque Area V4. *Neuron* **63**: 879–888.

Morcos, Ari S. and Christopher D. Harvey (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature Neuroscience* **19**.

Musall, Simon, Matthew T. Kaufman, Ashley L. Juavinett, Steven Gluf, and Anne K. Churchland (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience* **22**.

Namboodiri, Vijay Mohan K., Marco A. Huertas, Kevin J. Monk, Harel Z. Shouval, and Marshall G.Hussain Shuler (2015). Visually cued action timing in the primary visual cortex. *Neuron* **86**: 319–330.

Neske, Garrett T., Dennis Nestvogel, Paul J. Steffan, and David A. McCormick (2019). Distinct Waking States for Strong Evoked Responses in Primary Visual Cortex and Optimal Visual Detection Performance. *Journal of Neuroscience* **39**: 10044–10059.

Ni, A. M., D. A. Ruff, J. J. Alberts, J. Symmonds, and M. R. Cohen (2018). Learning and attention reveal a general relationship between population activity and behavior. *Science* **359**.

Niebergall, Robert, Paul S. Khayat, Stefan Treue, and Julio C. Martinez-Trujillo (2011). Multifocal attention filters targets from distracters within and beyond primate mt neurons' receptive field boundaries. *Neuron* **72**: 1067–1079.

Niell, C. M. and M. P. Stryker (2008). Highly Selective Receptive Fields in Mouse Visual Cortex. *Journal of Neuroscience* **28**: 7520–7536.

Niell, Cristopher M. and Michael P. Stryker (2010). Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. *Neuron* **65**: 472–479.

Nienborg, Hendrikje and Bruce G Cumming (2006). Macaque V2 Neurons, But Not V1 Neurons, Show Choice-Related Activity. *Journal of Neuroscience*.

— (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* **459**: 89–92.

Ohki, Kenichi, Sooyoung Chung, Yeang H. Ch'ng, Prakash Kara, and R. Clay Reid (2005). *Functional imaging with cellular resolution reveals precise microarchitecture in visual cortex.*

O'Keefe, John (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology* **51**: 78–109.

Olshausen, Bruno A. and David J. Field (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**: 607–609.

Ott, Torben, Paul Masset, and Adam Kepecs (2018). The neurobiology of confidence: From beliefs to neurons. *Cold Spring Harbor Symposia on Quantitative Biology* **83**: 9–16.

Pachitariu, Marius, Nicholas A Steinmetz, Shabnam N Kadir, Matteo Carandini, and Kenneth D Harris (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort. *Advances in Neural Information Processing Systems 29*. Ed. by D D Lee, M Sugiyama, U V Luxburg, I Guyon, and R Garnett. Curran Associates, Inc.: 4448–4456.

Parker, A. J. and W. T. Newsome (1998). SENSE AND THE SINGLE NEURON: Probing the Physiology of Perception. *Annual Review of Neuroscience* **21**: 227–277.

Peterson, Matthew F. and Miguel P. Eckstein (2012). Looking just below the eyes is optimal across face recognition tasks. *Proceedings of the National Academy of Sciences of the United States of America* **109**: E3314–E3323.

Petruno, Sarah K., Robert E. Clark, and Pamela Reinagel (2013). Evidence That Primary Visual Cortex Is Required for Image, Orientation, and Motion Discrimination by Rats. *PLoS ONE* **8**: e56543.

Poort, Jasper, Adil G. Khan, Marius Pachitariu, Abdellatif Nemri, Ivana Orsolic, Julija Krupic, Marius Bauza, Maneesh Sahani, Georg B. Keller, Thomas D. Mrsic-Flogel, and Sonja B. Hofer (2015). Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. *Neuron* **86**.

Rao, Rajesh P.N. and Dana H. Ballard (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* **2**: 79–87.

Raposo, David, Matthew T. Kaufman, and Anne K. Churchland (2014). A category-free neural population supports evolving demands during decision-making. *Nature Neuroscience* **17**: 1784–1792.

Raposo, David, John P. Sheppard, Paul R. Schrater, and Anne K. Churchland (2012). Multisensory decision-making in rats and humans. *Journal of Neuroscience* **32**: 3726–3735.

Reid, R. Clay and Jose Manuel Alonso (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* **378**: 281–284.

Reig, Ramon and Gilad Silberberg (2014). Multisensory Integration in the Mouse Striatum. *Neuron* **83**: 1200–1212.

Resulaj, Arbora, Sarah Ruediger, Shawn R. Olsen, and Massimo Scanziani (2018). First spikes in visual cortex enable perceptual discrimination. *eLife* **7**.

Reynolds, John H and Leonardo Chelazzi (2004). Attention modulation of visual processing. *Annu. Rev. Neurosci* **27**: 611–658.

Rigotti, Mattia, Omri Barak, Melissa R. Warden, Xiao Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**: 585–590.

Roelfsema, P R and H Spekreijse (2001). The representation of erroneously perceived stimuli in the primary visual cortex. *Neuron* **31**: 853–63.

Roelfsema, Pieter R and Floris P de Lange (2016). Early Visual Cortex as a Multiscale Cognitive Blackboard. *Annual review of vision science* **2**: 131–151.

Roitman, Jamie D. and Michael N. Shadlen (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience* **22**: 9475–9489.

Romo, Ranulfo and Emilio Salinas (2003). Cognitive neuroscience: Flutter Discrimination: Neural codes, perception, memory and decision making. *Nature Reviews Neuroscience* **4**: 203–218.

Rossi, Mark A., Volodya Y. Hayrapetyan, Benjamin Maimon, Krystal Mak, H. Shawn Je, and Henry H. Yin (2012). Prefrontal cortical mechanisms underlying delayed alternation in mice. *Journal of Neurophysiology* **108**: 1211–1222.

Saleem, Aman B., E. Mika Diamanti, Julien Fournier, Kenneth D. Harris, and Matteo Carandini (2018). Coherent encoding of subjective spatial position in visual cortex and hippocampus. *Nature* **562**.

Salzman, C Daniel, Chieko M Murasugi, Kenneth H Britten, and William T Newsome (1992). Microstimulation in Visual Area MT: Effects on Direction Discrimination Performance. *Journal of Neuroscience* **72**: 2331–2355.

Samejima, Kazuyuki, Yasumasa Ueda, Kenji Doya, and Minoru Kimura (2005). Representation of action-specific reward values in the striatum. *Science* **310**: 1337–1340.

Sanders, Joshua I and Adam Kepecs (2012). Choice ball: a response interface for two-choice psychometric discrimination in head-fixed mice. *Journal of Neurophysiology* **108**: 3416–3423.

Sanders, Joshua I. and Adam Kepecs (2014). A low-cost programmable pulse generator for physiology and behavior. *Frontiers in Neuroengineering* **7**: 1–8.

Schuett, Sven, Tobias Bonhoeffer, and Mark Hübener (2002). Mapping retinotopic structure in mouse visual cortex with optical imaging. *Journal of Neuroscience* **22**: 6549–6559.

Scott, Benjamin B., Carlos D. Brody, and David W. Tank (2013). Cellular Resolution Functional Imaging in Behaving Rats Using Voluntary Head Restraint. *Neuron* **80**: 371–384.

Scott, Benjamin B, Christine M Constantinople, Jeffrey C Erlich, David W Tank, and Carlos D Brody (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *eLife* **4**: 1–23.

Scott, Benjamin B., Christine M. Constantinople, Athena Akrami, Timothy D. Hanks, Carlos D. Brody, and David W. Tank (2017). Fronto-parietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales. *Neuron* **95**: 385–398.e5.

Shang, Congping, Zijun Chen, Aixue Liu, Yang Li, Jiajing Zhang, Baole Qu, Fei Yan, Yaning Zhang, Weixiu Liu, Zhihui Liu, Xiaofei Guo, Dapeng Li, Yi Wang, and Peng Cao (2018). Divergent midbrain circuits orchestrate escape and freezing responses to looming stimuli in mice. *Nature Communications* **9**.

Shiozaki, Hiroshi M., Seiji Tanabe, Takahiro Doi, and Ichiro Fujita (2012). Neural activity in cortical area V4 underlies fine disparity discrimination. *Journal of Neuroscience* **32**: 3830–3841.

Shuler, Marshall G. and Mark F. Bear (2006). Reward timing in the primary visual cortex. *Science* **311**: 1606–1609.

Sippy, Tanya, Damien Lapray, Sylvain Crochet, and Carl C.H. Petersen (2015). Cell-Type-Specific Sensorimotor Processing in Striatal Projection Neurons during Goal-Directed Behavior. *Neuron* **88**: 298–305.

Spitzer, Hedva, Robert Desimone, and Jeffrey Moran (1988). Increased attention enhances both behavioral and neuronal performance. *Science* **240**: 338–340.

Sriram, Balaji, Alberto Cruz-Martin, Lillian Li, Pamela Reinagel, and Anirvan Ghosh (2018). A Sparse Unreliable Distributed Code Underlies the Limits of Behavioral Discrimination. *bioRxiv*: 424713.

Stanisor, Liviu, Chris Van Der Togt, Cyriel M.A. Pennartz, and Pieter R. Roelfsema (2013). A unified selection signal for attention and reward in primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* **110**: 9136–9141.

Steinmetz, Nicholas A, Peter Zatka-Haas, Matteo Carandini, and Kenneth D Harris (2019). Distributed coding of choice, action and engagement across the mouse brain Brain-wide recording in visual behaviour. *Nature* **576**.

Stringer, Carsen, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D Harris (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**.

Toscani, Matteo, Matteo Valsecchi, and Karl R. Gegenfurtner (2013). Optimal sampling of visual information for lightness judgments. *Proceedings of the National Academy of Sciences of the United States of America* **110**: 11163–11168.

Trenholm, Stuart and Arjun Krishnaswamy (2020). An Annotated Journey through Modern Visual Neuroscience. *Journal of Neuroscience* **40**: 44–53.

Vinck, Martin, Renata Batista-Brito, Ulf Knoblich, and Jessica A. Cardin (2015). Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding. *Neuron* **86**: 740–754.

Wachowiak, Matt (2011). All in a Sniff: Olfaction as a Model for Active Sensing. *Neuron* **71**: 962–973.

Wallace, Damian J., David S. Greenberg, Juergen Sawinski, Stefanie Rulla, Giuseppe Notaro, and Jason N.D. Kerr (2013). Rats maintain an overhead binocular field at the expense of constant fusion. *Nature* **498**: 65–69.

Wang, Lupeng and Richard J. Krauzlis (2018). Visual Selective Attention in Mice. *Current Biology* **28**.

Wang, Quanxin and Andreas Burkhalter (2007). Area Map of Mouse Visual Cortex. *Journal of Comparative Neurology* **357**: 339–357.

Wimmer, Ralf D., L. Ian Schmitt, Thomas J. Davidson, Miho Nakajima, Karl Deisseroth, and Michael M. Halassa (2015). Thalamic control of sensory selection in divided attention. *Nature* **526**: 705–709.

Wolpert, Daniel M., Zoubin Ghahramani, and Michael I. Jordan (1995). An internal model for sensorimotor integration. *Science* **269**: 1880–1882.

Xin, Yu, Lin Zhong, Yuan Zhang, Taotao Zhou, Jingwei Pan, and Ning long Xu (2019). Sensory-to-Category Transformation via Dynamic Reorganization of Ensemble Structures in Mouse Auditory Cortex. *Neuron* **103**: 909–921.e6.

Xiong, Qiaojie, Petr Znamenskiy, and Anthony M. Zador (2015). Selective corticostriatal plasticity during acquisition of an auditory discrimination task. *Nature* **521**: 348–351.

Yamins, Daniel L.K. and James J. DiCarlo (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience* **19**.

Yang, Hongdian, Sung E. Kwon, Kyle S. Severson, and Daniel H. O'Connor (2015). Origins of choice-related activity in mouse somatosensory cortex. *Nature Neuroscience* **19**: 127–134.

Yang, Scott Cheng Hsin, Máté Lengyel, and Daniel M. Wolpert (2016). Active sensing in the categorization of visual patterns. *eLife* **5**.

Yarbus, Alfred L (1967). *Eye Movements and Vision*. Springer US. ISBN: 978-1-4899-5379-7.

Yilmaz, Melis and Markus Meister (2013). Rapid innate defensive responses of mice to looming visual stimuli. *Current Biology* **23**: 2011–5.

Zador, Anthony M. (2019). *A critique of pure learning and what artificial neural networks can learn from animal brains*.

Zatka-Haas, Peter, Nicholas A. Steinmetz, Matteo Carandini, and Kenneth D. Harris (2019). Distinct contributions of mouse cortical areas to visual discrimination. *bioRxiv*: 501627.

Zmarz, Pawel and Georg B. Keller (2016). Mismatch Receptive Fields in Mouse Visual Cortex. *Neuron*.

Znamenskiy, Petr and Anthony M. Zador (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature* **497**: 482–485.

"*...If you find this work difficult and wearisome to follow,*

*take pity on me,*

*for I have repeated these calculations seventy times...*"

Johannes Kepler