Learning enhances encoding of time and temporal surprise in primary sensory cortex

Rebecca Rabinovich

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2022

Abstract

Learning enhances encoding of time and temporal surprise in primary sensory cortex

Rebecca Rabinovich

Primary sensory cortex has long been believed to play a straightforward role in the initial processing of sensory information. Yet, the superficial layers of cortex overall are sparsely active, even during strong sensory stimulation; moreover, cortical activity is influenced by other modalities, task context, reward, and behavioral state. The experiments described in this thesis demonstrate that reinforcement learning dramatically alters representations among longitudinally imaged neurons in superficial layers of mouse primary somatosensory cortex. Cells were confirmed to be sparsely active in naïve animals; however, learning an object detection task recruited previously unresponsive neurons, enlarging the neuronal population sensitive to tactile stimuli. In contrast, cortical responses habituated, decreasing upon repeated exposure to unrewarded stimuli. In addition, after conditioning, the cell population as well as individual neurons better encoded the rewarded stimuli, as well as behavioral choice. Furthermore, in well-trained mice, the neuronal population encoded of the passage of time. We further found evidence that the temporal information was contained in sequences of cell activity, meaning that different cells in the population activated at different moments within the trial. This kind of time-keeping was not observed

in naïve animals, nor did it arise after repeated stimulus exposure. Finally, unexpected deviations in trial timing elicited even stronger responses than touch did. In conclusion, the superficial layers of sensory cortex exhibit a high degree of learning-dependent plasticity and are strongly modulated by non-sensory but behaviorally-relevant features, such as timing and surprise.

Table of Contents

List of Figures

Acknowledgments

I am deeply grateful to my research advisor, Randy Bruno, for his mentorship and support over the years. Randy has never failed to reassure me in stressful moments, offering guidance when everything felt doomed to fail, but also allowing me the space to formulate my own ideas. Thanks also to the rest of the Bruno lab for feedback, ideas, and camaraderie—your friendship has been central to my experience as a graduate student. In addition to the lab, the members of my thesis committee—Daphna Shohamy, Stefano Fusi, and Christoph Kellendonk—have been incredibly helpful with keeping me on track. Thanks also to my external examiner, Jerry Chen. I look forward to seeing you all at an in-person defense after two years of virtual meetings.

With respect to the specific research described in this thesis, I appreciate Dan Kato for teaching me how to image, for braving the pandemic to demonstrate surgery techniques, and for providing crucial feedback while I was writing up the project described in Chapter 2. I also must acknowledge Chris Rodgers for instruction and advice on previous versions of the behavioral apparatus and experiments (especially the task described in Chapter 3), for insightful comments on the manuscript (for the project in Chapter 2), and for first introducing me to coding in Python. Also, thanks to Ramon Nogueira and Stefano Fusi for advice on decoding analyses, Gordon Petty and Georgia Pierce for help and advice with pilot versions of whisking analyses, and Esther Greeman and Drew Baughman for help with histological sectioning.

My partner, Jack Bowler, has not only provided emotional support, but offered invaluable help with imaging and decoding analyses. Perhaps even more vital were the many conversations while walking the dog and over dinner in which he helped me brainstorm ideas—these discussions led me to ask the questions that turned into some of the most significant findings of my thesis.

Finally, not even my parents have been spared: Irina Rabinovich and Misha Rabinovich have listened to and critiqued countless practice talks, and have read and given feedback on almost every major piece of writing of my academic career—from grant proposals to drafts of my manuscript.

Dedication

This thesis is dedicated to my granddad, Semyon Rabinovich, who continuously worked to

foster my curiosity about the natural world during my childhood. Throughout his life, he

never lost his own curiosity, and was always eager to learn as well as to impart knowledge.

He would have loved to read this thesis.

Chapter 1 : Introduction

In order to survive in the world, humans and other animals must have adaptive interactions with their environment. Sense organs collect information about one's surroundings, and transmit this information to the brain, which must register these signals and interpret them, and then translate them into a behavioral output. Whether we see a movement in the leaves, hear thunder, or smell smoke, the information—visual, auditory, olfactory—must guide an appropriate behavior: depending on the sensory stimulus, we may freeze, seek shelter, or relocate.

The cerebral cortex, the six-layered structure thought to underlie mammalian cognition and sensory perception, has expanded over evolutionary history, presumably yielding richer sensory experiences and more complex behavioral responses to sensory stimuli. Among the sensory regions of the cortex, primary sensory cortex is traditionally believed to constitute an early stage in a sensory processing hierarchy. The simplest model of sensory information flow through primary sensory cortex can be summarized as follows: sensory information enters from the periphery, via sensory thalamus, and arrives at layer 4 of cortex; from here, signals continue to superficial layers (layers 2/3), which send information to other cortical areas; signals also project to deep layer 5, which relays information to subcortical brain structures. After the initial processing by primary sensory cortices, information is thought to be further manipulated by cortical regions "higher" in the hierarchy. This hierarchy of increasing complexity also implies that neural responses

should become sparser in higher-order areas: while primary cortical responses should be robust to a wider range of sensory stimuli (Avidan et al., 2002; Boynton et al., 1999), cells in higher-order areas are expected to be more selective.

In recent years, new information has changed our view of sensory processing, and these models have been updated. However, even decades ago, studies revealed that the response properties of the supposedly "downstream" deep layers do not necessarily depend on activity in the superficial layers: when the superficial layers of the visual cortex were inactivated, many cells in the deep layers remained unaffected (Schwark et al., 1986). More recently, Constantinople & Bruno (2013), as well as other authors since then (Egger et al., 2020; Pluta et al., 2019), identified thalamocortical inputs directly projecting to the deep layers of the somatosensory cortex (bypassing layer 4 and the superficial layers). Furthermore, the recorded responses in deep layer cells were unaffected by silencing of layer 4 (Constantinople and Bruno, 2013). Additional research confirmed that layer 4 and the superficial layers were not required for deep layer activity; in fact, layer 4 activity was found to *suppress* rather than drive layer 5 responses (Pluta et al., 2015). Moreover, while responses in layer 4 are indeed strongly driven by sensory stimuli, cells in superficial layers exhibit sparse activity (Barth and Poulet, 2012; Estebanez et al., 2012; Ramirez et al., 2014).

Finally, the presence of top-down inputs to primary sensory cortices suggests that these "early" sensory regions may be involved in more complex computations than originally believed: feedback connections from association cortices modulate activity in primary sensory regions, rendering responses in these areas more sensitive to expectation, attention, and motivation (Gilbert and Sigman, 2007; Poort et al., 2022). A balance

between top-down feedback projections and bottom-up feedforward inputs is required for adaptive sensation to be preserved; imbalances have been theorized to play a role in conditions like schizophrenia (Gilbert and Sigman, 2007; Teufel and Fletcher, 2020).

Indeed, evidence increasingly suggests a more multifaceted role of primary sensory cortex, in which the activity of any given primary sensory cortical region is influenced by a variety of factors, such as other modalities, task context, reward, and behavioral state (Brosch et al., 2011; Budinger et al., 2006; Lacefield et al., 2019; Mima et al., 1998; Pantoja et al., 2007; Pleger et al., 2008; Rodgers et al., 2021; Shuler and Bear, 2006; Weis et al., 2013; Zhang et al., 2020). In this chapter, I will examine the ways in which primary sensory cortical activity depends on elements other than bottom-up sensory inputs from the periphery, focusing first on the effects of learning and then on the influence of expectation (and violations of expectations). Then, I will propose that a function of primary sensory cortex is to use temporal processing to model the external environment and internal state of the animal—this model may then be used to predict outcomes and guide actions.

1.1   Primary sensory cortical activity is modulated by learning

Learning the relationship between a stimulus and an outcome is likely to affect primary sensory cortical responses, as evidenced by studies of multiple modalities, and across a variety of animal models. In some cases, cells in primary sensory cortical areas exhibit responses to reward, even in the absence of the sensory stimulus. For instance, the activity of some neurons in the deep layers of rat primary visual cortex (V1) not only

responded to visual stimuli, but also predicted the timing of subsequent reward delivery (Shuler and Bear, 2006). The reward-related signal was only observed after animals became proficient at the behavioral task—in naïve rats, cells only responded to visual stimuli, not reward; then, over the course of learning, a subset of neurons began to respond at the time of the reward as well. In fact, these cells responded at the expected time of reward delivery even when the reward itself was unexpectedly omitted. Reward-related activity has also been discovered in primary somatosensory cortex (S1): recordings in rat S1 revealed neuronal activity in anticipation of reward (Pantoja et al., 2007). An analogous human fMRI study found S1 to reactivate at the presentation of a reward following a tactile stimulus during a somatosensory discrimination task, though no tactile stimulus was present during this second peak of activity (Pleger et al., 2008); this reward-related signal, in turn, was found to correlate with performance on the behavioral task (Pleger et al., 2009, 2008). A calcium imaging study of mouse S1 also revealed a late peak of activity in both deep and superficial layers; the authors hypothesized that it was related to reward delivery rather than stimulus presentation (Lacefield et al., 2019). Specifically, delaying the reward delivery by a variable amount led to a corresponding delay of this late calcium signal. Finally, a study of monkey auditory cortex also identified various forms of reward-related activity: among different groups of cells, response magnitudes varied depending on reward expectations, reward size, and prediction error (Brosch et al., 2011).

Not only do primary sensory cortical cells respond to reward, but also their *stimulus-driven* responses may be modulated by learning: numerous studies have illustrated that learning leads to changes in cortical representations. Even three decades ago, early studies

in guinea pigs identified a degree of plasticity in primary sensory (auditory) cortex (Bakin and Weinberger, 1990). More recently, in superficial layers (layer 2/3) of mouse V1, cells were found to become more selective to a rewarded oriented grating stimulus after mice learned an association between the stimuli and reward (Henschke et al., 2020). In fact, across the neuronal population, cells' orientation preferences shifted toward the orientation of the rewarded stimulus (Henschke et al., 2020). Stimulus representations also stabilized as animals learned a behavioral task (Poort et al., 2015). However, it does seem that different characteristics of behavioral paradigms (such as whether the outcome is appetitive or aversive) yield different effects on primary cortical activity: Makino & Komiyama (2015) saw that layer 2/3 cells responded to stimuli even in naïve mice when those stimuli were paired with aversive outcomes. Nonetheless, they did still observe changes in the cells' activity over the course of learning, with neuronal responses peaking in anticipation of the aversive event as mice learned the avoidance task (Makino and Komiyama, 2015). Learning-related plasticity of stimulus-responses by primary sensory cortical cells is further supported by structural modifications observed in S1 during conditioning (Kuhlman et al., 2014): a rapid proliferation and stabilization of dendritic spines was detected on cells in superficial layers of S1 over the course of reinforcement learning. Interestingly, the rate of spine growth correlated with animals' behavioral performance, further solidifying the hypothesis that cortical plasticity really does occur as an effect of learning. These results match the observation that learning leads to an increase in magnitude of the responses of some groups of S1 cells (Chen et al., 2015).

On the other hand, Bale et al. (2020) found that learning did not have much direct effect on the sensory responses of S1: cells in superficial layers responded to the sensory features of the stimuli even before training. Yet, this study also reported that a subset of neurons seemed to gain representations of the stimulus-reward association over the course of training, such that neural responses in well-trained mice appeared to correspond to behavioral choice. Specifically, cells in expert animals responded to non-sensory aspects of the task (including animal's motor actions like licking) and predicted future decisions (Bale et al., 2020).

These changes in primary sensory cortical activity likely occur as a result of the learned stimulus-outcome associations, rather than some other aspect of the experiment such as repeated exposure to the stimulus. (Though, for evidence to the contrary, see Gavornik and Bear, 2014). In both primary auditory and visual cortices, responses to task-relevant or rewarded stimuli were shown to be modulated by learning, but passive stimulus exposure did not reproduce these changes (Henschke et al., 2020; Xin et al., 2019). Indeed, while passive repeated exposure to a stimulus leads to *habituation* of sensory responses (Henschke et al., 2020; Kato et al., 2015), that same stimulus can again begin to elicit strong responses if it becomes behaviorally-relevant (Kato et al., 2015).

1.2   Expectation, predictability, and attention

Stimulus and reward responses that emerge only in well-trained animals also suggest that expectations influence primary sensory cortical activity. In fact, several

authors have proposed predictive coding as the critical operation performed by the cortex (den Ouden et al., 2009; Keller and Mrsic-Flogel, 2018).

Indeed, a study in human auditory cortex study that identified reward-related signals found this activity to depend not on the delivery of the reward itself, but rather on the correct *prediction* of the presence or absence of reward (Weis et al., 2013). Similarly, as described in the previous section, reward-related activity in rat V1 only presented in well-trained animals that fully expected reward delivery (Shuler and Bear, 2006)—although, notably, whether the expected reward was actually presented or omitted did not impact the cells' responses. Both of these studies considered the predictability of reward, but other experiments showed that the predictability of sensory stimuli also affects neuronal responses. Rendering stimuli predictable via attentional cueing enhanced sensory responses to those stimuli (Doherty et al., 2005). Conversely, unexpected sensory events have been found to decrease cortical activity: though a cessation of tactile flow led to both up- and down-modulation of activity in the primary somatosensory cortex of mice, the majority of cells displayed weaker activity in response to the perturbation (Ayaz et al., 2019)—this effect was observed in both superficial and deep layers of SI, though the down-modulation was stronger in superficial layers. Finally, familiarity with stimuli has been showed to augment cortical responses, which were diminished when expectations were disturbed (Gavornik and Bear, 2014).

However, there exist even more cases in which primary sensory cortical cells responded more strongly when expectations were violated rather than confirmed. Reward-related activity in the primate primary auditory cortex (AI) was sensitive to mismatch

between expected and presented rewards, and some cells specifically tracked animals'

behavioral errors by selectively responding to unexpectedly omitted rewards on error trials

(Brosch et al., 2011). With respect to sensory (rather than reward) mismatch signals,

neurons in mouse primary visual cortex that respond to optic flow showed stronger

responses when the visual stimulus was unexpectedly stopped (Keller et al., 2012). In fact,

violations of expectations can drive cortical activity even in the absence of any sensory

stimulus: if an expected visual stimulus is *omitted*, cells in V1 have been found to exhibit

strong responses (Fiser et al., 2016). Meanwhile, predictable stimuli yielded much lower

responses in human V1, compared to unexpected stimuli (Alink et al., 2010), consistent

with the well-documented phenomenon of repetition suppression (Henson, 2003;

Summerfield et al., 2008).

"*Repetition suppression*" refers to the dampened sensory responses to repeated

stimuli compared to "oddball" stimuli that pop out and recruit relatively more neural

resources. This phenomenon has been commonly explained as either a "low-level" process

like adaptation (Grill-Spector et al., 2006), or as an effect of expectation and predictability

(Summerfield et al., 2008). The latter hypothesis is supported by cases in which a stimulus

may be expected but not repeated (*e.g.*, if it is predicted by another stimulus or if it is

presented repeatedly but alternated with another stimulus)—a related phenomenon which

may be called "*expectation suppression*" (Richter et al., 2018).

An alternative explanation to many of these findings has been proposed as well:

rather than being a result of predictive coding processes, the responses to an expected

stimuli may be suppressed due to decreased attention allocated to that stimulus (Alink and

Blank, 2021). Whether the responses of primary sensory cortical regions to unexpected events are enhanced due to the high degree of prediction error associated with those events or due to attentional modulation remains an open question.

1.3   Role of primary sensory cortices

If primary sensory cortices do not simply process and relay sensory information from the periphery to so-called "higher order" association areas, then what does primary sensory cortex do? One potential role of cortex could be in acquiring the ability to perform complex feature recognition. Such a theory could explain why primary sensory cortex appears to be dispensable for simple tasks (Hong et al., 2018). I will revisit this possibility in Chapter 3, in which I describe a feature discrimination task intended to probe the involvement of primary sensory cortex in these types of computations.

An alternative, but not mutually exclusive, role of cortex could be to use predictive coding and temporal processing to build and update a model of an animal's external environment and internal state. In the following section, I elaborate on this idea, reviewing evidence of time coding in various regions of the brain, and finishing with a discussion of primary sensory cortical temporal processing.

1.4   Temporal processing around the brain

The ability to keep track of time and predict the timing of events is crucial for an animal's survival. For instance, timing important for predictions, as well as the ability to understand causal relationships between events or actions (Woods et al., 2014): we perceive an earlier event as the "cause", and a subsequent event as an "effect". Indeed, this temporal perception is required for proper prediction and interpretation of the effects of one's own actions. Delusions of control seen in schizophrenia may be an example of this system gone awry: control of one's actions by an outside entity may be a logical explanation if one perceives one's action to precede intent. Temporal processing disturbances have in fact been proposed to play a role in the symptoms of schizophrenia (Martin et al., 2014).

The importance of time-keeping has led to a long-time search for a neural "clock". Various candidate regions have been proposed to enable an animal to keep track of time— from cerebellum (Ashe and Bushara, 2014) to hippocampus and the neighboring lateral entorhinal cortex (Tsao et al., 2018) to striatum (Bakhurin et al., 2017; Mello et al., 2015). Yet sensory cortex has not been considered. I do *not* propose that sensory cortex constitutes such a clock; on the contrary, I argue that time-keeping is not localized to any one region, but rather constitutes a critical computation that is widely distributed across the brain, including primary sensory cortices.

Consider the diverse evidence for temporal processing around the brain: "*time cells*" in hippocampus respond sequentially over the course of a trial, and can link events across a delay (MacDonald et al., 2011; Pastalkova et al., 2008; Umbach et al., 2020). Meanwhile,

in the dorsolateral striatum, cells have also been found to track time (Jin et al., 2009; Toso et al., 2021b; Zhou et al., 2020), notably, even when timing is not informative for the task (Toso et al., 2021b). Time could also be decoded in certain cases in the amygdala and some association cortices (Cueva et al., 2020). However, this last study only found time to be encoded in these regions when temporal information was task-relevant, suggesting that some brain regions only use timing information to model the world when currently necessary, while others may encode temporal information more generally.

It is important to note that time encoding is unlikely to be an epiphenomenon: it does not inevitably arise from the dynamics of brain activity. There are cases in which time cannot be decoded from activity even in brain regions like hippocampus, where one would expect cells to encode temporal information. For instance, Ahmed, Priestley et al. (2020) were unable to decode time from hippocampal area CA1 during a trace fear conditioning task, likely due to the high degree of variability in neuronal responses across trials (Ahmed et al., 2020). Thus, certain conditions must be met for time encoding to occur.

Temporal processing in primary sensory cortex

While no studies have directly investigated temporal encoding in primary sensory cortices, there exists some indirect evidence of time processing in sensory cortex. For instance, primary somatosensory cortex has been found to be sensitive to stimuli characterized by temporal sequences (whisker deflections separated by variable intervals), suggesting the presence of temporal integration in S1 (Pitas et al., 2017). In addition,

temporally patterned SI activity has been shown to contain sensory information (Arabzadeh et al., 2006). A related finding in primate SI uncovered "memory cells" (Bodner et al., 2005), which displayed precise patterns of activity during a delay, and were therefore hypothesized to contribute to working memory.

Primary sensory cortices are also involved in temporal expectations: in well-trained animals, V1 responses predict the *timing* of an expected reward (Shuler and Bear, 2006) or aversive outcome (Makino and Komiyama, 2015). Meanwhile, in humans, temporal expectations yielded V1 activity in areas that matched the retinotopic location of a stimulus expected to appear at a certain time, even in the absence of the stimulus (Bueti et al., 2010). These observations imply the existence in primary visual cortex of a temporal model of sequences of events, experiences, and expectations.

Indeed, V1 neurons have been shown to "learn" the temporal order of sequences of events, such that they respond more strongly to an *expected* sequence than to a novel sequence. Gavornik and Bear (2014) reported that, as mice become familiarized with a given sequence of visual stimuli (gratings of various orientations) over the course of multiple days, the visual responses to that sequence become progressively stronger with training. On the other hand, control animals that are presented with random sequences, rather than the same sequence each time, exhibit less enhancement in V1 responses. This discrepancy is notable because the constituent visual stimuli of the sequences are the identical; only the order of presentation is randomized. Eventually, mice in the experimental cohort come to fully expect the specific sequence on which they are trained. Subsequently, if the constituent elements of the sequence are either reordered or adjusted

in temporal duration, the visual responses decrease: even though the individual stimuli are the same, adjusting the order or timing of the events is sufficient to render the composite sequence "novel". Meanwhile, control animals display no differences between V1 responses to the various sequences, as no one sequence is "novel" nor "familiar" (Gavornik and Bear, 2014).

Surprisingly, Gavornik and Bear (2014) did not use reward to enable expectations: repeated exposure was sufficient for the sequence to become expected and to augment cortical responses. Another study, however, *did* pair reward with a target stimulus to help animals form expectations about the timing of that stimulus presentation (Jaramillo and Zador, 2011). This study also found an effect of temporal expectations, wherein a context of anticipation enhanced stimulus-evoked primary sensory cortical responses: neurons in rat A1 became more sensitive to auditory stimuli that occurred *around the time* of an expected target (rewarded) stimulus than to stimuli that occurred at moments farther away from the target in time. In this experiment, a series of identical tones was punctuated by rewarded "target" tones of a different frequency, which signaled to the rats that a reward was to be delivered—the exact pitch of the target tone informed animals about the location of the reward. This target tone could arrive at different moments ("early" or "late") in the sequence of non-rewarded "baseline" tones, with alternating blocks of trials, such that the target arrived "early" with a high probability in one block of trials, but "late" in the next. This experimental design enabled animals to learn to predict the timing of the target's arrival. As the moment of presentation of an expected rewarded tone drew closer, cortical activity driven by the non-rewarded tones of other frequencies was enhanced: tones immediately

preceding the target elicited stronger A1 responses than did tones earlier in the trial. Interestingly, this modulation of cortical activity by expectation only affected the stimulus-driven A1 responses, and not spontaneous activity. Furthermore, the expectation-dependent enhancement of cortical responses was stimulus-specific, and was more evident in neurons with frequency preferences that matched the stimuli (Jaramillo and Zador, 2011).

Finally, sensory aspects of a stimulus can influence time perception in both humans and other animals (Herbst et al., 2013; Reinartz et al., 2021; Toso et al., 2021b, 2021a; Tse et al., 2004; Xuan et al., 2007), a phenomenon that suggests that timing and sensation are likely linked. The relationship between temporal and sensory processing is further substantiated by experiments that reveal bidirectional perceptual changes in temporal duration (as well as in reported sensation of the stimulus intensity) following optogenetic manipulation of primary sensory cortex (Reinartz et al., 2021). In this study, two separate groups of rats were trained to discriminate the intensities or durations of tactile stimuli. The authors then optogenetically stimulated or inhibited S1, and found that upregulating the cortex led to judgements of longer duration, while inhibiting S1 caused animals to perceive the stimulus to be of shorter duration. A parallel study in humans (Salvioni et al., 2013) found that interfering with V1 activity using transcranial magnetic stimulation (TMS) affected participants' ability to discriminate interval durations. Finally, in the cutaneous rabbit illusion (Geldard and Sherrick, 1972; Figure 1.1), which may depend on temporal integration of the sequential sensory inputs, a signature of the illusory percept has been identified in human S1 (Blankenburg et al., 2006). In this illusion, the forearm is repeatedly

and rapidly tapped at discrete areas—a simple variation of the illusion involves three taps at the wrist followed by two taps at the inner elbow. These taps are perceived as progressing ("hopping") up the arm, rather than being localized to just the wrist and elbow. The aforementioned human neuroimaging study (Blankenburg et al., 2006) found that the somatotopic activation in SI during the illusion is identical to that of a sequence of taps actually progressing up the arm.
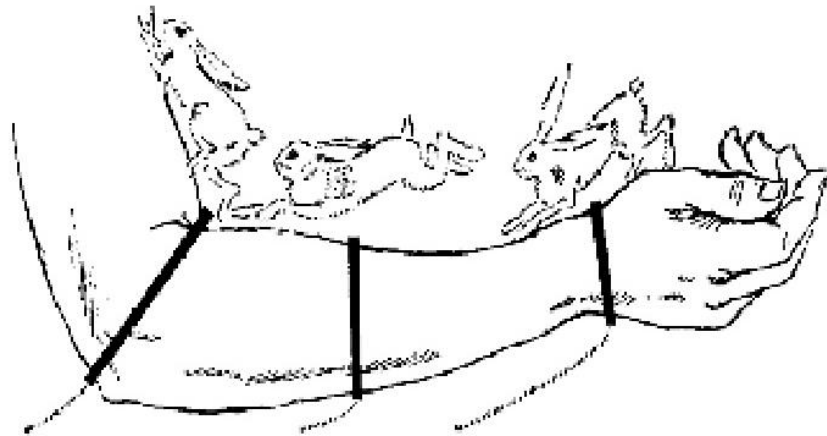
Figure 1.1 The cutaneous rabbit illusion
The cutaneous rabbit illusion (Geldard and Sherrick, 1972) involves a series of taps on discrete areas of the forearm, and is felt as a sequence of taps "hopping" up the arm.

Together, these studies suggest that primary sensory cortices do contribute to temporal processing. Still, there are notable gaps in our understanding of their role in these processes. No direct investigation has been done to determine whether any primary sensory cortical region actually keeps track of time. Furthermore, while it seems likely that

expectations (including temporal expectations) modulate activity in primary sensory cortex, the nature of this effect is unclear. The studies described in this section seem to imply that cortical cells respond more strongly when the timing of sensory stimuli or rewards is *expected* (Gavornik and Bear, 2014; Jaramillo and Zador, 2011; Shuler and Bear, 2006*)*; yet, earlier, I mentioned other studies that reported increased cortical responses to *unexpected* sensory occurrences (Alink et al., 2010; Keller et al., 2012). However, these latter studies did not focus on the temporal properties of the surprising events. The research I conducted in this thesis aimed to resolve some of these discrepancies and gaps.

1.5   Rodent barrel cortex as an experimental model system

The studies summarized above span a wide range of modalities and animal models. This thesis utilizes the rodent barrel cortex—the portion of primary somatosensory cortex that receives tactile sensory input from an animal's whiskers—as a model system to investigate primary sensory cortical involvement in temporal processing and how it develops over learning. The motivation behind choosing the mouse barrel cortex as a model system is two-fold, relating to (1) choice of animal model, and (2) choice of cortical area:

1) A mouse model allows for invasive experimental procedures, including visualization of individual cell activity in vivo.

2) Given the choice of the mouse model, the whisker system is an appropriate choice of modality as it is one which mice frequently rely on in nature: since mice spend their time in dark underground tunnels, their whiskers provide a large

proportion of their sensory information about the environment. However, my

assumption is that my findings would apply equally to other modalities, and I

hope that my results can be replicated in other primary sensory cortical regions.

The barrel cortex is a particularly well-studied cortical region, with strict

somatotopic mapping between an animal's whiskers and columnar structures in the cortex.

These columns, or "barrels", lend this cortical region its name. Figure 1.2 shows a tangential

slice of the barrel cortex, with barrels visible as darker areas.

Figure 1.2: The barrel field
This 50um layer 4 tangential slice of barrel cortex shows the barrel field. The darker regions are the barrel columns. Each barrel receives sensory input from a single whisker.

1.6   Overview

In the following chapter, I will describe a series of experiments that probe activity

in barrel cortex: specifically, I will illustrate how learning a behavioral task (even in a very

17

simple object detection paradigm) alters cortical representations in a number of ways—modulating sensory responses as well as the encoding of behavioral choice, time, and temporal expectations. Then, in Chapter 3, I will consider a more complex behavioral paradigm, provide preliminary behavioral data, and outline the challenges that a more complex experimental design brings. Finally, I will discuss how my findings relate to theories of temporal processing in cortex and beyond, and propose a direction for future research into the neural correlates of time perception.

Chapter 2 : Effects of learning on primary sensory cortical activity and temporal coding

This Chapter is adapted from Rabinovich, Kato, & Bruno (under review; [preprint on BioRxiv](#)). I performed the imaging experiments and analyses; Dan Kato performed some of the surgeries, and contributed animals to the experiments; he also provided valuable feedback on the manuscript under review.

2.1   Introduction

The established view of cortical processing assumes that primary sensory cortex performs basic tasks at an early stage in a sensory processing workflow, with more complex computations occurring downstream (for review see Grill-Spector and Malach, 2004). However, there is mounting evidence that more "associative" processing occurs in primary sensory cortices and that activity in these regions is influenced by factors beyond simple sensation of a single modality (Brosch et al., 2011; Budinger et al., 2006; Lacefield et al., 2019; Mima et al., 1998; Pantoja et al., 2007; Pleger et al., 2008; Rodgers et al., 2021; Shuler and Bear, 2006; Weis et al., 2013; Zhang et al., 2020). Additionally, the canonical view implies robust stimulus responses by primary sensory cortex (Avidan et al., 2002; Boynton et al., 1999); on the contrary, primary sensory cortical cells, specifically superficial layer

neurons, often exhibit low levels of activity even during strong sensory stimulation (Barth and Poulet, 2012; Estebanez et al., 2012; Ramirez et al., 2014).

A large volume of research now suggests that sensory-driven responses in primary cortical areas are modulated by learning and stimulus-reward associations (Bakin and Weinberger, 1990; Brosch et al., 2011; Chen et al., 2015; David et al., 2012; Makino and Komiyama, 2015; Pantoja et al., 2007; Poort et al., 2015; Shuler and Bear, 2006; Weis et al., 2013). In the superficial layers (layer 2/3) of primary visual cortex (V1), initially quiet cells begin responding to behaviorally relevant stimuli that are paired with reward (Henschke et al., 2020), with representations becoming more selective and stable over the course of learning (Poort et al., 2015). Meanwhile, structural changes have been found to occur in layer 2/3 of primary somatosensory cortex (Kuhlman et al., 2014), where certain groups of cells respond more strongly to touch after learning (Chen et al., 2015). Yet, the degree to which learning-related plasticity occurs in primary somatosensory cortex (S1), and how it manifests, remains unclear: in some cases, stimulus representations have been found to remain stable (Kim et al., 2020), with little plasticity (Peron et al., 2015).

Mismatch between expectation and sensation has also been proposed to influence activity in primary sensory cortex (Ayaz et al., 2019; Keller et al., 2012). In fact, prediction, mismatch, and predictive coding in general have been theorized to play a crucial role in cortical processing, and may drive learning-related plasticity (den Ouden et al., 2009; Keller and Mrsic-Flogel, 2018). For instance, Keller et al. (2012) showed that disturbances in optic flow, such as the sudden cessation of visual motion, led to increased V1 activity (Keller et al., 2012). However, an alternative mechanism has been proposed for this finding

(Muzzu and Saleem, 2021), whereby V1 responses increase whenever visual motion slows due to cells' velocity preferences, regardless of direction of motion, or whether the change violated the animals' expectations (but see Keller et al., 2012 and Zmarz and Keller, 2016, who did not observe altered V1 activity during passive viewing of the stimulus. Meanwhile, perturbations in expected tactile flow mainly result in decreased neuronal activity in barrel cortex (the whisker-related portion of S1) (Ayaz et al., 2019). Thus, despite the increasing consensus regarding the importance of predictive coding (Fletcher and Frith, 2009; Keller and Mrsic-Flogel, 2018), the nature of mismatch signals in primary sensory cortices is still poorly understood.

The seemingly disparate observations across these studies of primary sensory cortex may be interconnected if primary sensory cortices construct a model over the course of learning, placing stimuli, rewards, and other events into a temporal context. Various brain regions have been already implicated in time processing. For instance, compelling evidence of "time cells" has been demonstrated in the hippocampus. These cells specifically respond at given moments in a trial and have been proposed to contribute to working and episodic memory by linking events across a delay (MacDonald et al., 2011; Pastalkova et al., 2008) and by representing their sequential order (Umbach et al., 2020). Similar sequence-driven coding schemes may be at play in striatum (Toso et al., 2021b; Zhou et al., 2020). Some degree of time encoding was also identified in the amygdala and several association cortical areas (Cueva et al., 2020). In sensory cortex, barrel cortical neurons have been shown to be sensitive to temporally patterned stimuli (Pitas et al., 2017); conversely, temporally-precise patterns of neuronal activity have been found to contain non-temporal somatosensory

information (Arabzadeh et al., 2006). Similarly, putative "memory cells" in primate somatosensory cortex exhibit temporal patterns of activity during a tactile working memory task (Bodner et al., 2005). In fact, multiple of the aforementioned studies that found learning-related influences on V1 activity noted a timing component to the changes they observed (Makino and Komiyama, 2015; Shuler and Bear, 2006).

The current study asks: to what degree does learning impact the activity of neurons in superficial layers of primary somatosensory cortex? What aspects of the behavioral paradigm do these cells encode? What kind of mismatch signals are present in S1? We used 2-photon calcium imaging to measure cortical activity during learning of a simple Pavlovian detection task. We found that learning a stimulus-reward association engages previously unresponsive cells in a longitudinally-tracked population, leading to an enhanced representation of tactile stimuli. In addition, learning rendered cells able to encode *non-sensory* aspects of the behavioral paradigm and animals' experience, including animals' choice, the temporal progression of events, and the expectation of event timing.

2.2   Results

To investigate learning-related effects on cortical activity, we used 2-photon calcium imaging to monitor the activity of GCaMP6f-expressing neurons in superficial layers (layer 2/3) of S1, while mice learned a Pavlovian whisker-based object-detection task (Figure 2.1). In our behavioral paradigm, stimuli were presented to water-restricted mice via a rotating wheel (Figure 2.1a, top). On rewarded trials, a flat surface rotated into the whisker field,

stopped and remained stationary for two seconds, and then rotated away, after which a water reward was presented (Figure 2.1a, bottom; see 2.4   Methods). On unrewarded trials, the wheel rotated to an empty position in which no surface contacted the whiskers, and no water was given. Thirsty mice licked at the water port in anticipation of the reward, and we quantified the level of anticipatory licking as a measure of learning. Initially, mice licked indiscriminately on both stimulus-present and stimulus-absent trials (Figure 2.1b, left), but after learning the association between stimulus and reward, they licked preferentially during stimulus-present trials, in anticipation of reward (Figure 2.1b, right; n = 10 mice).

The Pavlovian nature of the behavioral paradigm meant that an animal's actions did not impact the outcome of the trial (whether or not a water droplet emerged from the lick-port). Nonetheless, to quantify learning, we labelled mouse responses as follows: anticipatory licking on a stimulus-present trial was categorized as a "hit", suppression of licking on a stimulus-absent trial was a "correct rejection", licking on a stimulus-absent trial was a "false alarm", and lack of licking on a stimulus-present trial was a "miss". The former two response types were considered "correct", and the latter two were "incorrect". By these definitions, mice began to perform above chance in as few as three sessions, but we continued training them until their performance reached 70% (Figure 2.1c); in some cases, animals' performance exceeded 90%. Performance on this task fell to chance when whiskers were trimmed off (in 5 out of 5 whisker-trimmed mice).

Prior to training and 2-photon imaging (Figure 2.1d), we used intrinsic signal imaging to locate the barrel cortex, where we would later record neuronal activity. After

23

recording, we confirmed the imaging location by using the 2-photon laser to create a small lesion at the imaging site, and, with the help of vasculature patterns, identifying the corresponding location relative to the layer 4 barrels in post hoc histology (Figure 2.1e).
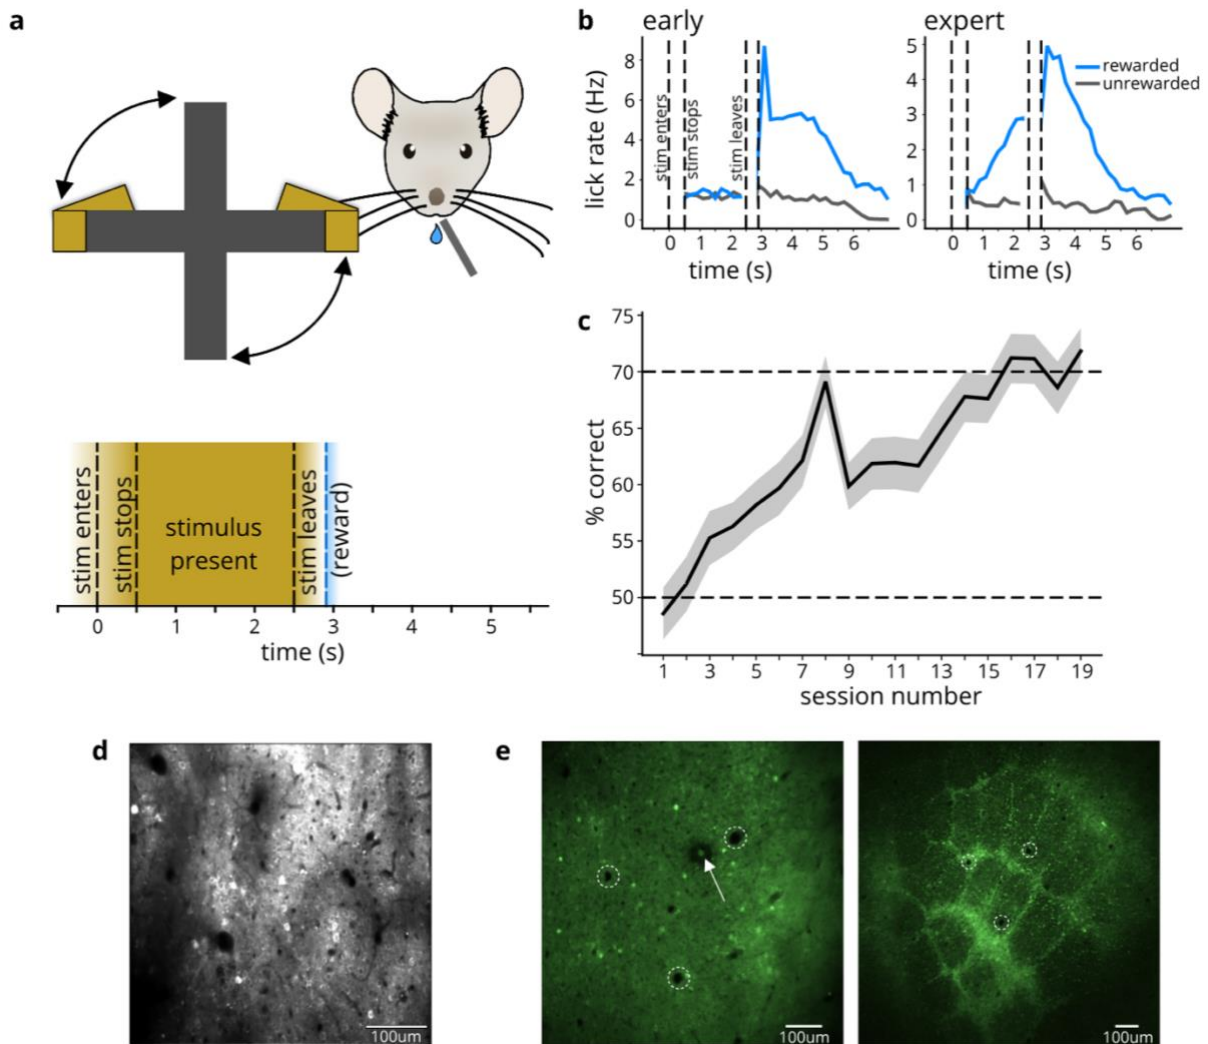
Figure 2.1: L2/3 calcium imaging in barrel cortex while mice learn whisker-based object detection task
a) Behavior schematic (top) and timecourse (bottom): rotating wheel brings either object or empty space into mouse whisker field. Stimulus "enters" the whisker field and then "stops". The now stationary stimulus is present for two seconds before rotating away (45 degrees over the course of ~400ms, such that the gap between wheel arms is in front of the animal's face). On object-present trials, water drop is

given following the stimulus. Anticipatory licks are counted during the "stimulus present" interval.

b) Anticipatory licking histograms (n = 10 mice) for early (left) and expert (right) days, quantifying the rate of licking for stimulus-present rewarded (blue) trials and stimulus-absent unrewarded (black) trials. Mice initially lick equally for object-present (rewarded) and object-absent (unrewarded) conditions, but learn to lick preferentially for the object-present condition with training.

c) Learning curve (n=10 mice). Shading corresponds to 95% confidence interval.

d) Example field of view for calcium imaging during learning. Calcium activity was measured in GCaMP6f-expressing neurons in layer 2/3 of barrel cortex.

e) Histology shows location of imaging site. Left: L2/3 tangential slice, showing lesion location (marked with arrow). Right: L4 tangential slice, showing corresponding area of barrel map. Dashed circles indicate blood vessels used to match location in barrel map.

Conditioning, but not repeated stimulus exposure, enhances object representation

Across the neuronal population, cellular activity was variable: neurons responded at different times within the trial, both in naïve and expert sessions (see example mouse in Figure 2.3a, top). For most mice, average responses to the stimulus increased in amplitude as mice became proficient at the detection task: on early days, the average population activity was relatively flat; after conditioning, the activity peaked in response to the onset and offset of the stimulus (Figure 2.3a, bottom; same example mouse). In addition to analyzing the overall change in fluorescence, we extracted the times of calcium transients (see 2.4   Methods; Figure 2.2). All subsequent analyses are based on the time or rate of these calcium transients, unless otherwise noted. We quantified the neuronal population response for each mouse on early and expert sessions as the mean number of calcium

transients across cells. Population responses to object arrival increased from early to expert days (p = 0.001, paired t-test, n = 10 mice; Figure 2.3b).
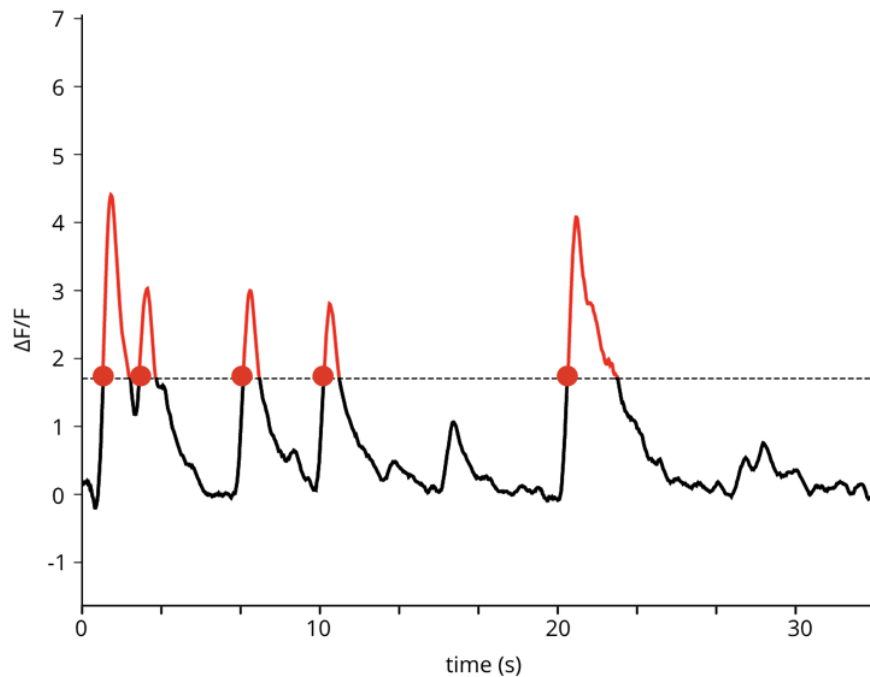
Figure 2.2: Detecting calcium transients from fluorescence. Threshold for each cell is set as 2 standard deviations above the median ΔF/F for that cell. The time when the trace first crosses the threshold is considered to be the transient onset time. Suprathreshold periods are colored red; transient onsets are marked with red circles.

Averaging signals across all cells obscures the variability in cell responses within the population, so we next examined individual cell activity. The responses of individual cells (n_{early} = 1163 cells, n_{expert} = 1112 cells) could be categorized into several groups based on the timing of their calcium transients (Figure 2.3c; see 2.4 Methods for classification criteria). Some cells responded specifically at the onset of the stimulus ("on" cells), while others responded at the stimulus offset ("off" cells); another group responded at both onset and offset ("on-off" cells); finally, a small subset only responded late in the trial, following

reward presentation ("reward" cells). The "on," "off," and "on-off" cells dominated layer 2/3 SI activity, as can be seen in the average rate of calcium transients across all imaged cells: transients occur most frequently at stimulus onset and offset, as well as, to a lesser extent, the period in between (Figure 2.3d). This effect was more pronounced in expert mice, consistent with the increased population of these cells.

In addition to these four response profiles, many cells displayed no obvious response patterns ("none" cells): in fact, in naïve mice, the majority of cells were in this category (Figure 2.3e, left). We refer to these neurons as unresponsive to contrast them with cells that exhibited immediate responses at the time of object arrival, departure, or reward delivery; in reality, this category may include cells that were active but did not meet our classification criteria for the other cell groups. As mice learned, the proportion of these unclassified "none" cells decreased ($p < 0.001$, Z approximation to a binomial) and the proportion of stimulus-responsive cells increased ($p < 0.001$ for "on" cells and $p < 0.001$ for "off" cells): in expert mice, most cells were stimulus-responsive (Figure 2.3e, right).

Was this substantial increase in the population of cells responsive to object arrival a genuine effect of conditioning, or could mere repeated exposure to stimuli be driving the increases in cells' responses? We explored this issue using a new cohort of animals—mice that were exposed to the same behavioral paradigm with respect to stimulus presentation, but were not water-restricted and never received rewards (n = 7 mice). Comparing the change in proportion of cells of each type as mice progressed through the protocols revealed that the presence of reward was in fact crucial for the increase in stimulus-responsive cells. Repeated exposure produced no increase in the percentage of responsive

cells; rather, we saw a *decrease* in the proportion of these cells (p < 0.001 for "on" cells; p = 0.03 for "off" cells, with unresponsive "none" cells growing to an even larger population (p < 0.001) after repeated exposure to the stimuli, an effect that was the *reverse* of that seen in conditioned mice (Figure 2.3f). These findings suggest that two opposing adaptation mechanisms exist in layer 2/3: irrelevant stimuli shrink the population of responsive cells while behaviorally important stimuli enlarge it.
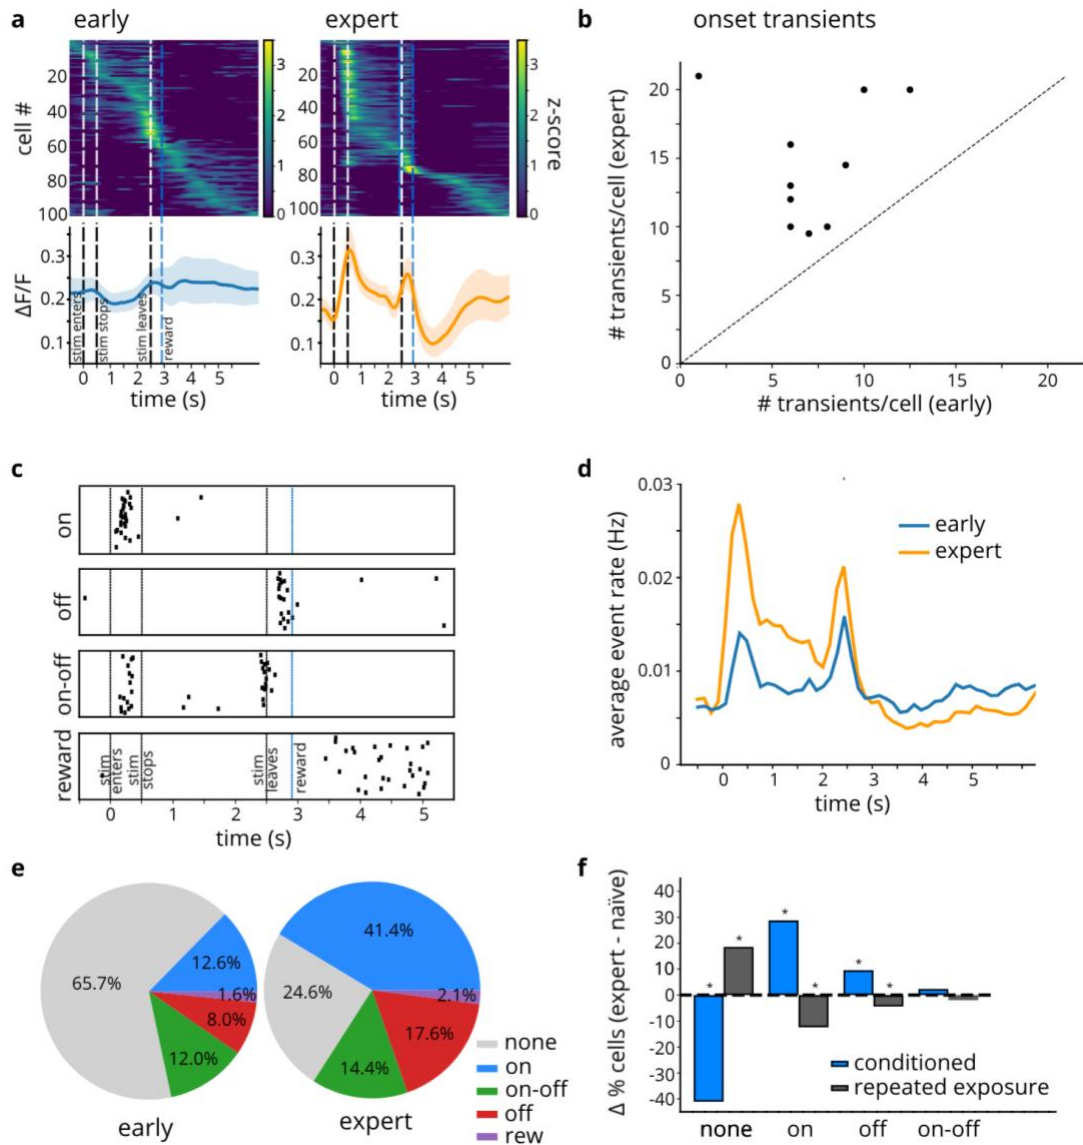
Figure 2.3: Behavioral training, but not stimulus exposure, increases proportion of stimulus-responsive cells (n = 10 mice)

a) Cells respond at different times in trial. Top: heat maps from 100 cells in a naïve session (left) and 104 cells in an expert session (right) in an example mouse; Bottom: average across cells for this mouse, for the same early (left) and expert (right) session.

b) Average number of transients per cell for early and expert sessions (p = 0.001, paired t-test). Each point corresponds to 1 mouse.

c) Raster plots of calcium transients for 4 example cells show diverse patterns of activity.

d) Transient event rate in naïve (blue) and expert (orange) animals. Rate of transients was calculated across all cells, normalized by number of trials and number of cells.

e) Proportions of different cell types in the population, on early vs expert days for conditioned mice: on: $p < 0.001$; off: $p < 0.001$; on-off: $p = 0.1$; rew: $p = 0.4$; none: $p < 0.001$ (based on Z-test approximation to the binomial; sample sizes: $n_{early} = 1163$ cells; $n_{expert} = 1112$ cells).

f) Change in the proportion of cells of each type (expert - naïve), for conditioned mice that experience reward-pairing (n=10) and for mice that only experience repeated exposure to the stimulus without reward (n=7). For repeated exposure mice: on: $p < 0.001$; off: $p = 0.03$; on-off: $p = 0.22$; none: $p < 0.001$. Data for conditioned mice same as in (e).

Learning switches response category of longitudinally tracked neurons

As mice learned the behavioral task above, a neuronal population initially dominated by unresponsive cells transformed into a population mainly comprising stimulus-responsive cells. However, this analysis was unable to reveal the dynamics of individual cells: what degree of fluctuation or stability did individual cells exhibit across consecutive days and over the course of learning? To address this question, we next set out to determine how *individual* cell responses shift over the course of conditioning.

To this end, we longitudinally tracked individual neurons across training sessions. To identify the same cells across two sessions, we warped the imaging field of view of one day to align with that of another day by applying an affine transformation to the time average of one imaging session; we then applied the same transformation to each region-of-interest (ROI) and located matching ROIs across sessions (Figure 2.4a; see 2.4 Methods).
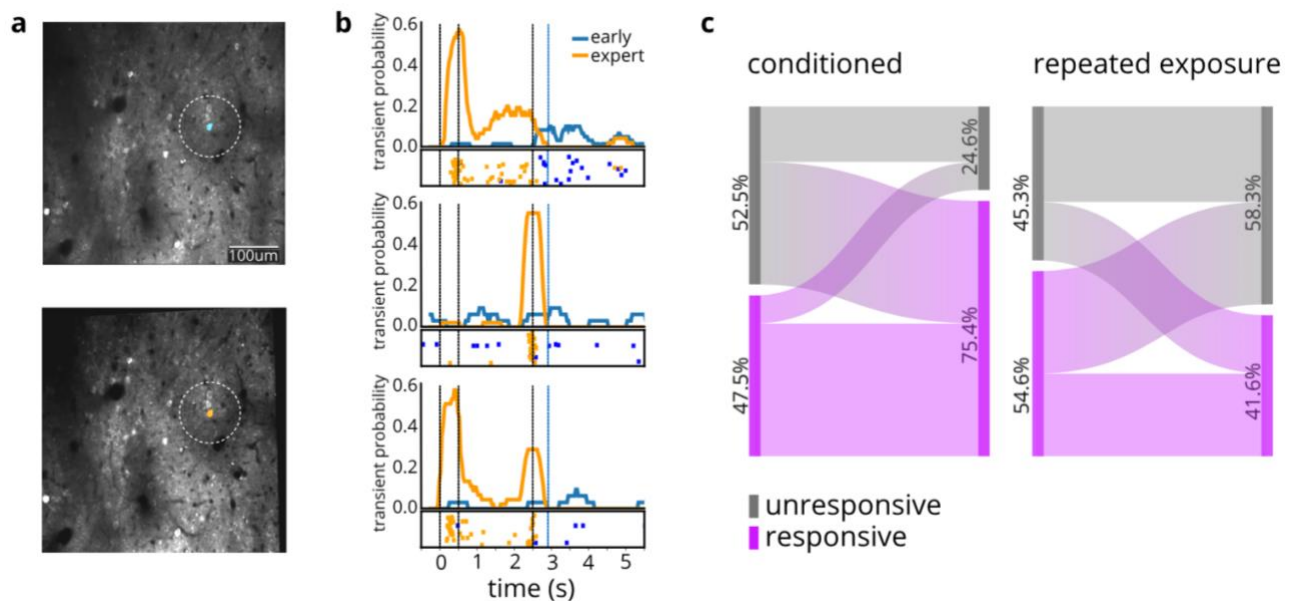
Figure 2.4: **Longitudinally tracked cells' responses exhibit opposing processes of enhancement and habituating in response to conditioning and repeated exposure, respectively.**
a) Field of view of example mouse on naïve (top) and expert (bottom) days. A single tracked cell is highlighted.
b) Three example cells' activity on early (blue) vs expert (orange) days. For each cell: PSTH (top) shows probability of calcium transients throughout the timecourse of a "hit" trial; raster plot (bottom) shows all calcium transients for all "hit" trials in the session.
c) Responsive and unresponsive cells longitudinally tracked across conditioning (left) or repeated exposure without reward (right). Percent cells for responsive and unresponsive cells was calculated for each mouse and then averaged across mice, to ensure that each mouse's contribution was weighted equally.

As mice progressed from naïve to expert levels of behavioral performance, longitudinally tracked cells' activity dramatically increased (see examples in Figure 2.4b). Averaged across mice, the majority (69%) of previously unresponsive cells began responding to the stimuli. Meanwhile, cells that were already responsive to the stimulus mostly remained responsive (83%), and only a small fraction (17%) became unresponsive (Figure 2.4c, left). For the repeated-exposure mice, on the other hand, tracked cells became *less* responsive to stimuli, with more than half (55%) of original responsive cells losing responsiveness, and most unresponsive cells (62%) remaining unresponsive (Figure 2.4c, right). The magnitudes of these changes were similar if we considered cells pooled across mice (Figure 2.5) rather than averaged.

However, in mice that had experienced reward-pairing, longitudinal cell tracking revealed that cells maintained a "memory" of the task and were relatively stable across consecutive days: 81% of cells were consistently within the same responsive/unresponsive category across consecutive expert sessions. Thus, over the course of learning, individual tracked cells' representations become biased toward reinforced stimuli, but only in conditioned mice. Later, the activity of these same cells stabilized, reflecting the animals' gained understanding of task rules.
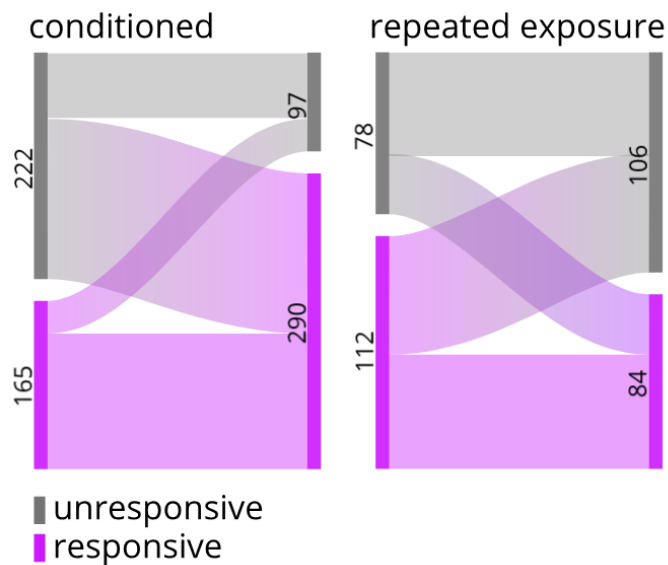
Figure 2.5: Longitudinally tracked cells become more responsive after conditioning, but not repeated exposure.
Same as Figure 2.4c, except cells from all mice are pooled together (rather than averaged across mice).

Training enhances decodability of stimulus and choice

Learning rendered cells not only more responsive, but also more predictive of trial type and the animal's behavior. We trained a support vector machine with a linear kernel to decode the stimulus (the presence or absence of the object) or choice (whether or not the mouse displayed anticipatory licking) from calcium transients in the neuronal population (Figure 2.6). On average, decoding accuracy for both stimulus (Figure 2.6a; paired t-test, p = 0.007) and choice (Figure 2.6b; paired t-test, p < 0.001) improved along with animals' performance.

The above analysis used all imaged neurons, including the unresponsive "none" cells. Non-classically responsive cells have previously been shown to be informative in auditory cortex (Insanally et al., 2019). Interestingly, our unresponsive cells also

contributed significantly to the decoders. Unresponsive cells had smaller contributions (decoder weights) in the stimulus classifier than did cells of other response profiles (Mann-Whitney U-test, p = 0.002) but had similar weights for choice decoding (Mann-Whitney U-test, p = 0.32). Yet even for the stimulus classification, decoding just from "none" cells yielded above chance performance on expert days (76.8% decoder performance; t-test, p < 0.001).
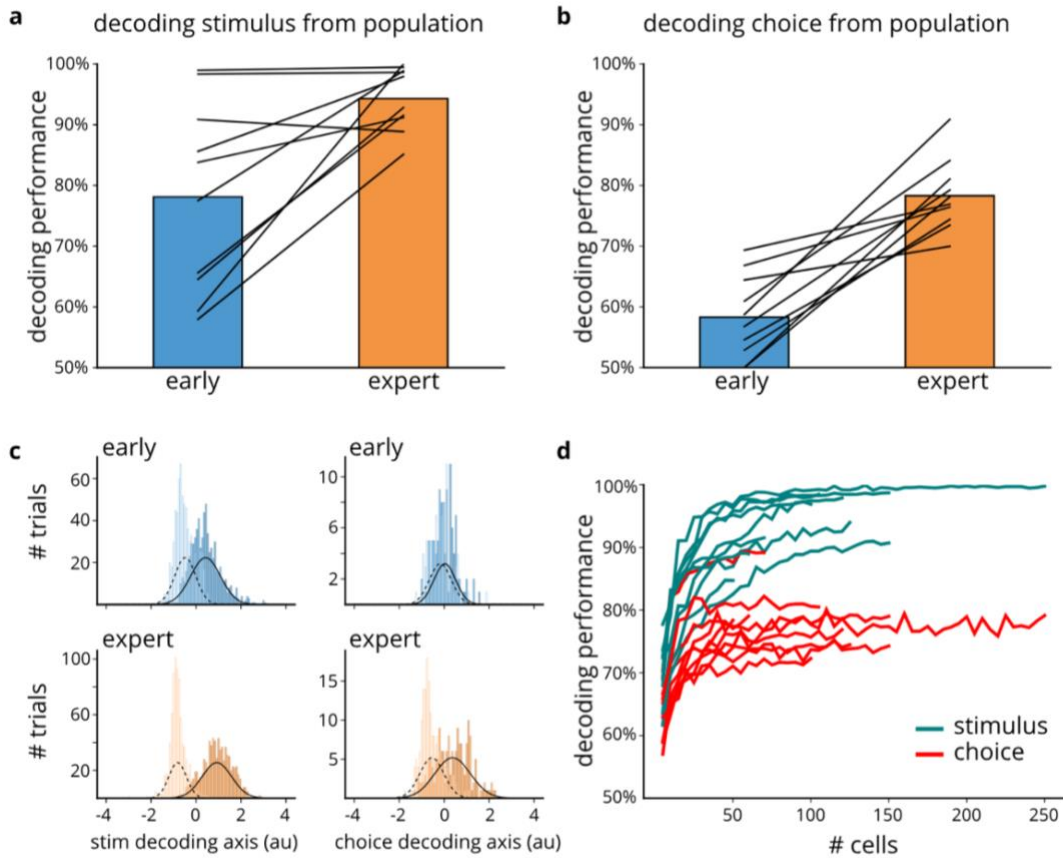
Figure 2.6: Training enhances decodability of stimulus and choice, from population activity.
a) Decoding stimulus from cell population on early (blue) and expert (orange) days (p=0.007). Black lines represent individual mice.
b) Decoding choice from cell population on early (blue) and expert (orange) days (p<0.001).
c) Separability of trial types when decoding from the cell population.
Left: stimulus-present vs stimulus-absent trials for early (top) and expert (bottom) days. (filled bars: stimulus present; empty bars: stimulus absent).
Right: trials with vs without anticipatory licks for early (top) and expert (bottom) days. (filled bars: lick; empty bars: no lick).
d) Decoding performance on expert days for stimulus (turquoise) and choice (red) using increasing number of cells. Curves correspond to individual mice.

Furthermore, even though the decoder could discriminate the trial type classes above chance even in naïve mice, the trial types became more separable on expert days compared to early-training days (Figure 2.6c): projecting neuronal activity across the coding axis (determined by the classifier weights) revealed that the trial type classes diverged on expert days, and therefore that the decoder was more "confident" in classifying the trial types (Figure 2.7a,b; average earth mover's distance between stimulus classes increased from 0.91 to 1.75; Wilcoxon signed rank test, p = 0.005; average earth mover's distance between choice classes increased from 0.18 to 1.02, p = 0.005). Note that for mice in the repeated exposure cohort, the ability to decode the stimulus from neuronal activity did not improve with continued stimulus presentation (Figure 2.7c), nor did the stimulus classes diverge (Figure 2.7d).
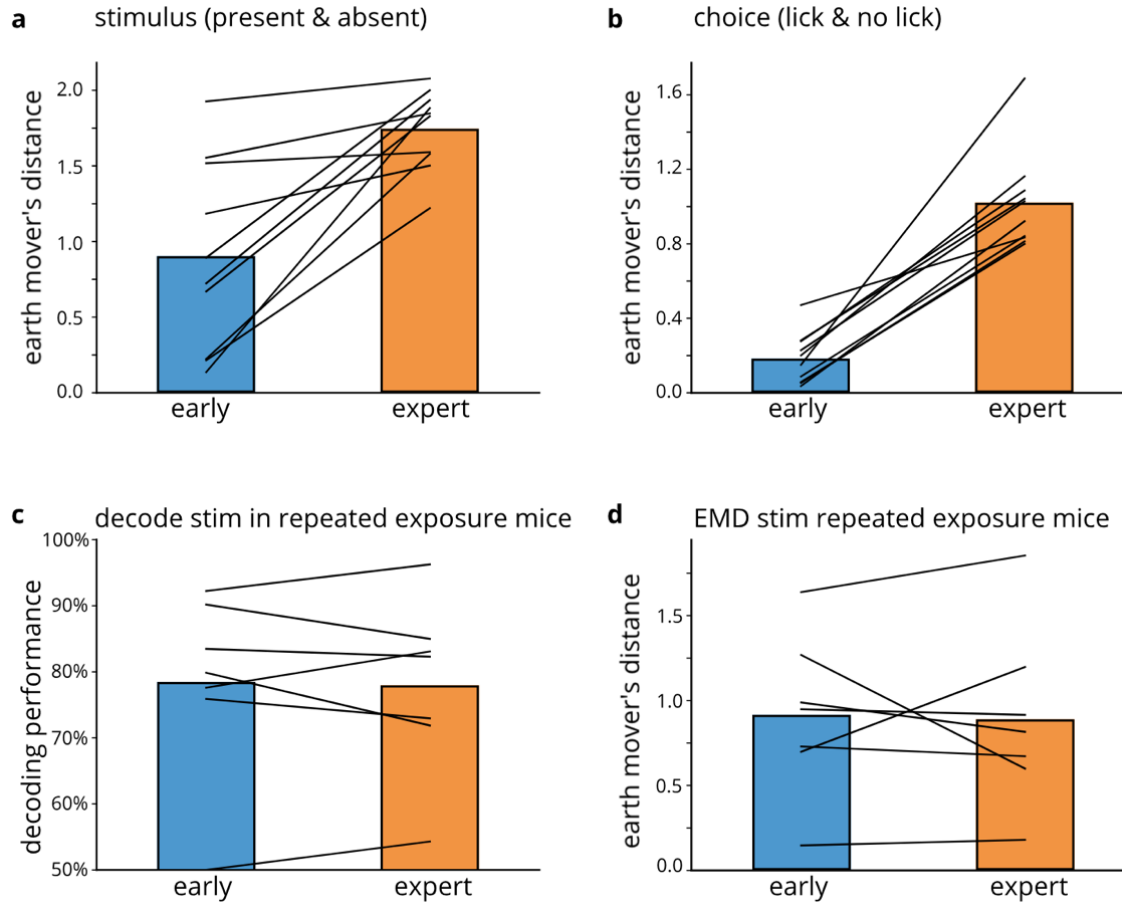
Figure 2.7: neural representations of trial types diverge with learning, but not following repeated exposure.
a) Earth mover's distance between stimulus classes (present & absent) on early and expert days. Black lines indicate individual mice; bars are the average (p = 0.005).
b) Same as (a) but for choice classes (lick & no lick) (p = 0.005).
c) Decoding stimulus from cell population of repeated exposure mice on early (blue) and late (orange) days (p = 0.925).
d) Same as (a) but for repeated exposure mice.

Since the total number of cells was not identical across different imaging sessions (ranging between 50-260 cells; mean = 116 cells), a potential confound could arise if decoding performance scales with population size. Comparing decoding accuracy using differently sized subsets of the same neuronal population relieved this concern by showing

that, while the decoding performance did initially increase with the number of cells used, it rapidly plateaued at sizes smaller than that of our typical imaging datasets (Figure 2.6d). In fact, decoding from just 50 cells was above chance (t-test, $p < 0.001$ for both stimulus and choice)—on average 93% for stimulus and 78% for choice. Thus, the different population sizes across sessions cannot explain the learning-dependent improvement in decodability of stimulus and choice.

So far, we have demonstrated that, at least on expert sessions, we were able to decode both stimulus and choice. However, these two measures become increasingly correlated as mouse performance improves: while mice never reach perfect performance, we expect stimulus and choice to be substantially correlated on expert days, when mice mainly choose to lick on stimulus-present trials and inhibit their response on stimulus-absent trials. Therefore, a problem arises: on expert days, cells' activity may truly predict only stimulus *or* choice, and our ability to decode both measures might stem from the correlation between them. To disambiguate our claim from this alternative explanation, we took advantage of the fact that mice do make errors even while performing relatively well, and we re-ran the decoding analysis using trial-balancing (Rodgers et al., 2021), whereby each trial type was weighted inversely to the frequency with which it occurred; in other words, rare trial types (miss and false alarm trials) were weighted more strongly than frequent trial types (hit and correct rejection trials). On expert days, when stimulus and choice are correlated, trial-balanced decoding performance for both measures was lower compared to decoding performance using an unbalanced approach (Figure 2.8a,b; paired t-test, stimulus $p = 0.03$, choice $p < 0.001$), but both remained significantly above chance

(Figure 2.8c; t-test, stimulus p < 0.001, choice p < 0.001). Consequently, while some degree of decoder performance is attributable to stimulus-choice correlation, the ability to decode one variable is not simply due to its correlation with the other.
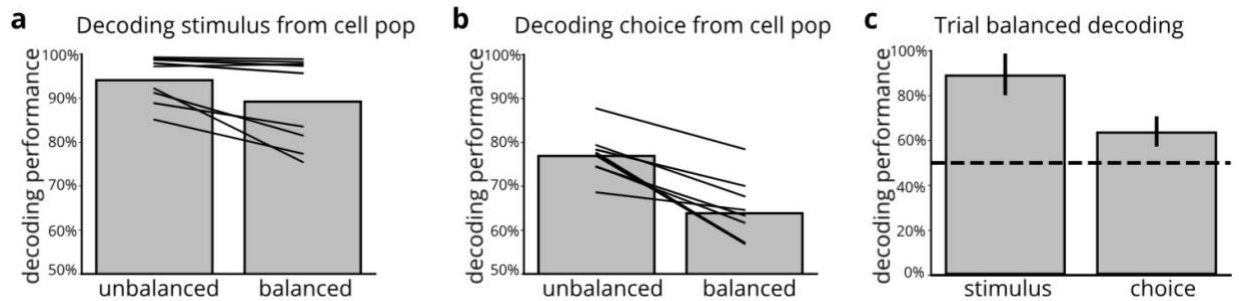
Figure 2.8: Trial-balanced decoding performance is lower than unbalanced approach, but still above chance.
a) trial-balanced stimulus decoding performance (right) compared to unbalanced (left) (p=0.03).
b) same as (a) but for choice (p = 0.0002).
c) 95% confidence interval for trial-balanced decoding (for stimulus and choice) above chance.

Stimulus and choice could also be decoded from individual cells, albeit to a lesser extent than from the entire population. On expert days, the decoder could predict trial type and the animal's response with greater accuracy from a larger percentage of single cells than on early days (Figure 2.9a). When compared to shuffled data, the stimulus and choice could be decoded above chance from more cells on expert days than on naïve days (Figure 2.9b; paired t-test, stimulus p = 0.005; choice p < 0.001). Finally, for individual, longitudinally tracked cells, decoding performance improved as the mouse learned the task (Figure 2.9c; paired t-test, stimulus p < 0.001; choice p < 0.001). Thus, as mouse behavioral

performance improved, individual cells, as well as population activity, became more predictive of trial type and mouse response type.

a decoding stimulus from single cells | decoding choice from single cells

b decoding stimulus above chance | decoding choice above chance

c decoding stimulus from tracked cells | decoding choice from tracked cells
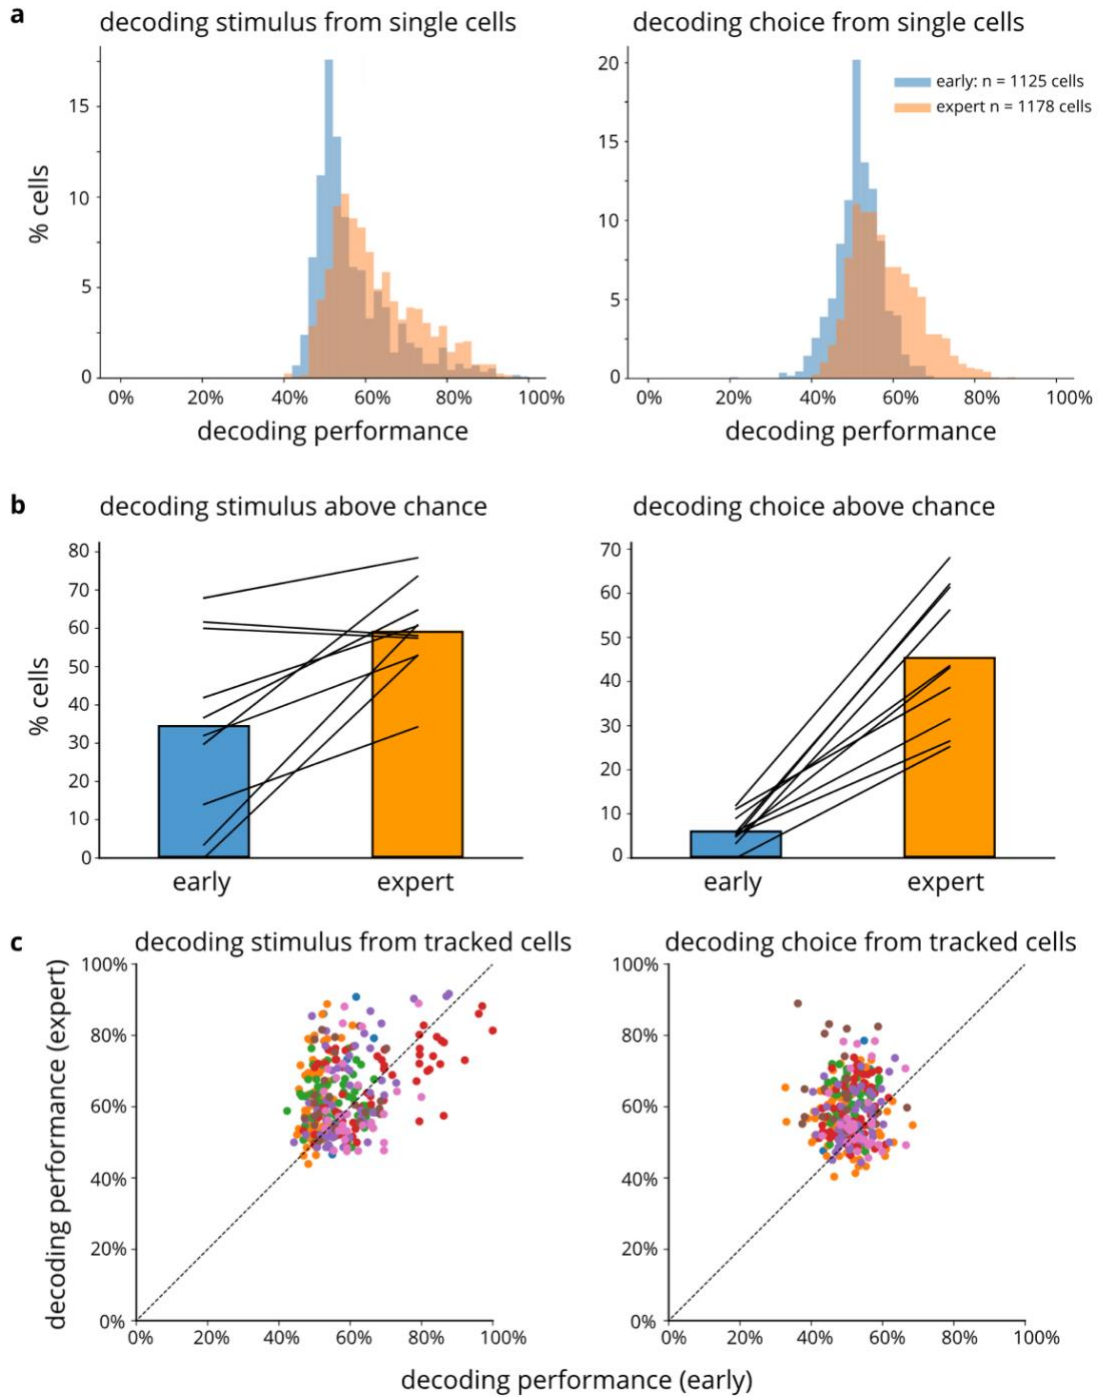
Figure 2.9: Training enhances decodability of stimulus and choice, from activity of individual cells.
a) Histograms of the percent of cells from which can decode stimulus (left) and choice (right) at a given accuracy, on naïve (blue) and expert (orange) days.
b) Percent of cells from which can decode stimulus (left) and choice (right) above chance, on naïve (blue) and expert (orange) days (stimulus p = 0.005; choice p < 0.001, paired t-test).
c) Decoding performance for stimulus (left) and choice (right), from individual longitudinally tracked cells, on naive vs expert days (stimulus p < 0.001; choice p < 0.001). Colors correspond to cells from different mice.

Temporal properties of neural representations

We next examined the temporal dimension of the neuronal representations by decoding stimulus and choice at various timepoints throughout the trial. We split the trials into time bins, with each bin containing 10 frames (1/3 of a second), and used the population activity at each time bin to decode stimulus and choice, yielding a decoding time course (Figure 2.10a). The number of time bins at which the decoder could predict stimulus and choice above chance increased as mice learned the task (Figure 2.10b; Wilcoxon signed rank test, stimulus p = 0.008, choice p = 0.005). Moreover, the first time bin from which trial type could be predicted shifted earlier over the course of learning (Figure 2.10c; Wilcoxon signed rank test, p = 0.03), suggesting that cells' activity encodes information about the stimulus (or future response) earlier in the trial as mice begin to comprehend the task demands and the relevance of the stimulus.
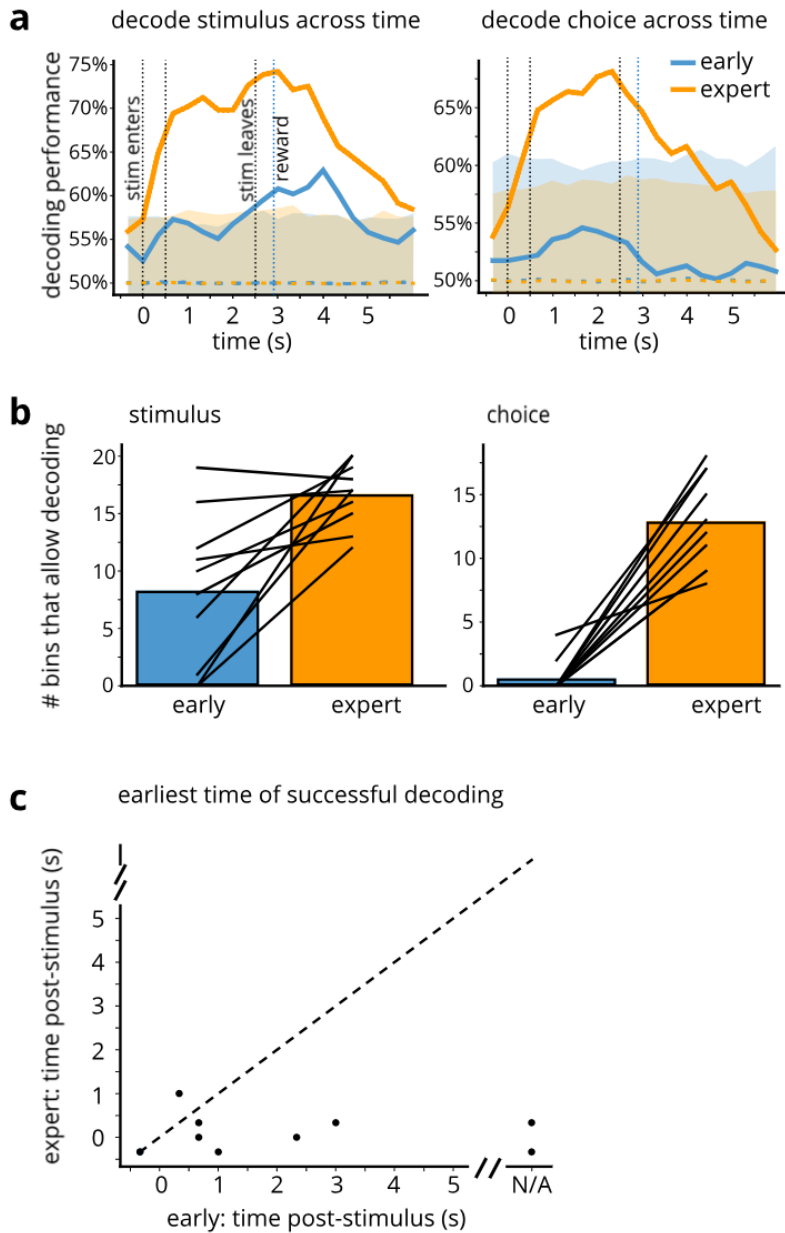
Figure 2.10: Temporal dynamics of neuronal representations.

a) Timecourse illustrating decoding performance, averaged across mice, for stimulus (left) and choice (right) from population of cells, for real neuronal data and for shuffled data. Blue: early day; orange: expert day. Dashed lines indicate mean of shuffled data; shaded area shows 99% confidence interval for shuffled data (n = 10 mice).

b) Number of time bins from which stimulus (left) and choice (right) can be decoded above chance (significantly different from shuffle), on naïve (blue) and expert (orange) days: above-chance decoding performance at more time bins on the trained day (stimulus: p = 0.008; choice: p = 0.005, Wilcoxon signed rank test; n = 10 mice).

c) First time in trial when stimulus can be decoded: can decode earlier in the trial after training (p = 0.03, Wilcoxon signed rank test). Each dot is one mouse.

Population activity encodes time

The above finding implies a temporal aspect to the cells' information coding, but the information encoded still only concerns trial parameters. Consequently, we asked: does neuronal activity also contain information about the *progression* of time? Specifically, could we predict the position of a given time bin within the temporal sequence from the population activity at that time? Indeed, in expert mice, we were able to decode time from calcium transients, not only at moments when a trial event was occurring (such as the stimulus arriving or leaving, or reward being delivered), but also during intervals when little was changing, such as the two-second period during which the stimulus was stationary in front of the mouse's face (Figure 2.11a, right). In naïve mice (Figure 2.11a, left) and in the repeated exposure cohort (Figure 2.12), the neuronal activity was less predictive of time bin identity, indicating that the temporal information encoded by the cell population arises due to learning. Analyses using smaller and larger bin sizes yielded qualitatively similar results.

We noticed that training on "hit" trials allowed the decoder to predict time on "hit" trials and "miss" trials, for which the stimulus was present, but not on "correct reject" trials, for which the stimulus was absent. We were unable to decode time on "correct reject" trials even when we trained the decoder on that trial type (Figure 2.11a). The inability to predict time on stimulus-absent trials could be due to the lack of tactile signals or lack of animal attention for that trial type.

41
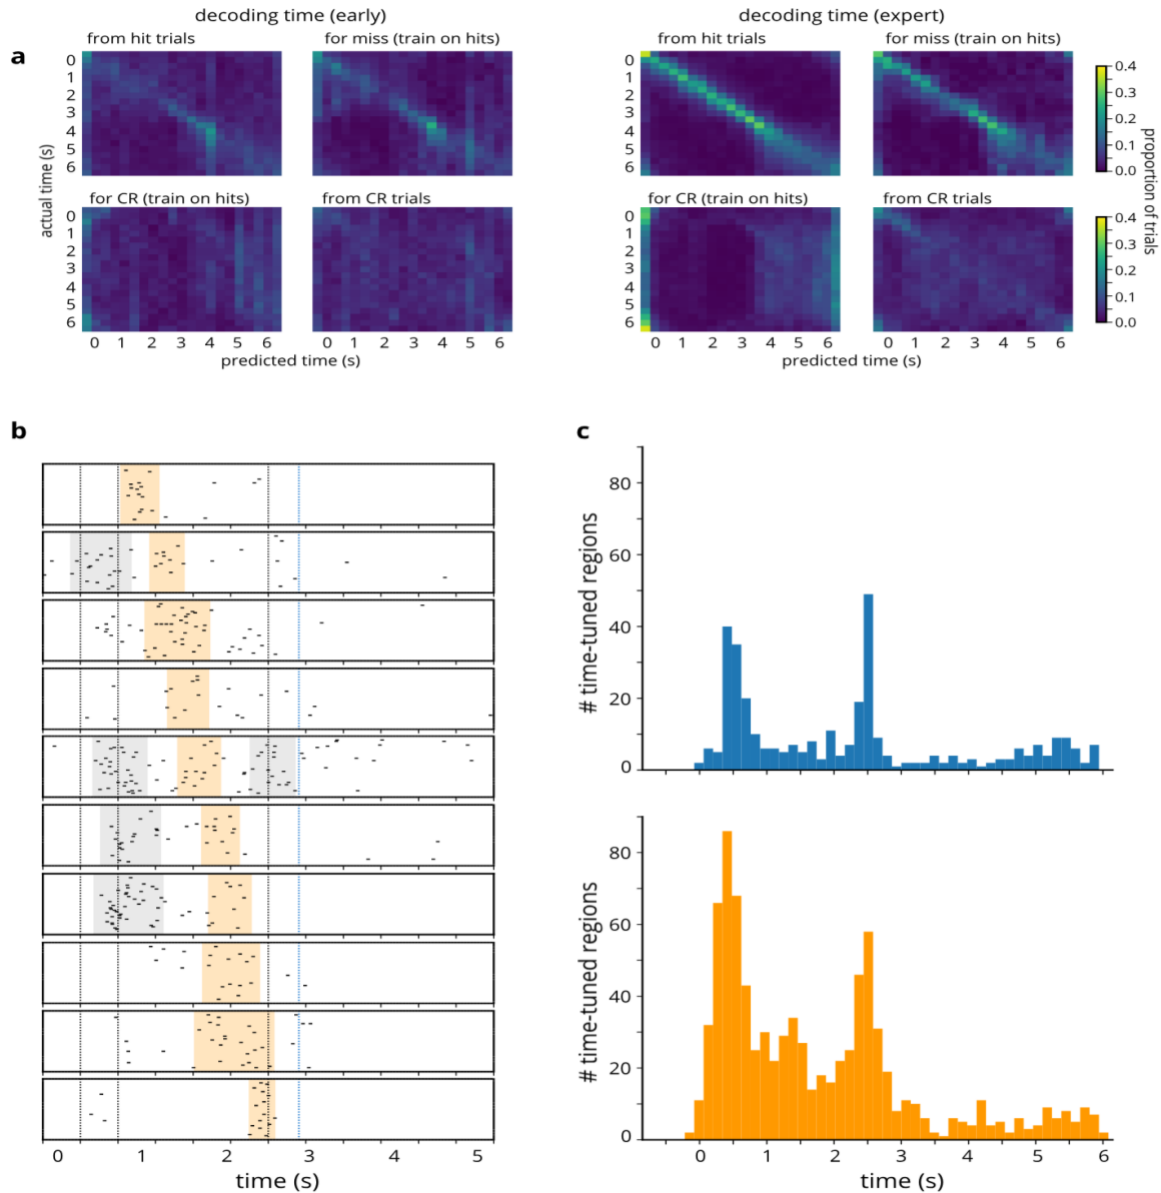
Figure 2.11: Cell population encodes time with high precision.

a) Decoding time from cell population, early (left) and after training (right). Columns indicate predicted time; rows indicate actual time. Colors show proportion of trials on which time y is predicted to be time x. Results were averaged across mice (n = 10 mice). Bin size = 10 frames. Sum across row = 1.

b) Cells are tuned to time intervals that span the duration of the trial. Shaded areas correspond to time-tuned regions. Some regions are highlighted (orange) to illustrate how these regions tile the trial timecourse.

c) Histograms indicating the total number of time-tuned regions in naïve (blue, top) and expert (orange, bottom) mice.

42
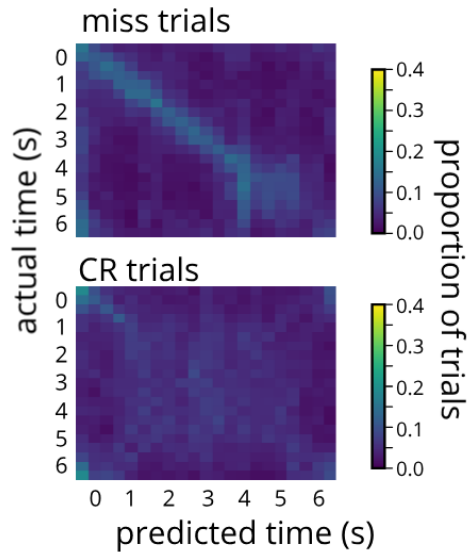
Figure 2.12: Poor time encoding in repeated exposure mice (n = 7).

Importantly, time keeping is not inevitable. Decoding time with high precision *on a trial-by-trial basis* from calcium transients in expert mice means that learning stabilizes previously unstructured trajectories in neuronal state space. Time decoding would be unsuccessful if these trajectories were inconsistent across trials (Ahmed, Priestley et al., 2020) or if they exhibited fixed-point dynamics (Cueva et al., 2020).

Having ascertained that the neuronal population encodes time with remarkable temporal resolution, we naturally asked how: *What coding scheme does barrel cortex use to track time?* Potential coding schemes include precise time-varying activity patterns of individual cells, stabilization of time constants of ramping activity, or response sequences across the neuronal population akin to those of hippocampal time cells. We found that neuronal activity can be described by the last option: cells are active during temporally-constrained periods that tile the duration of the trial (Figure 2.11b; see 2.4 Methods). Across the population of cells (Figure 2.11c), these time-tuned regions cluster around the

stimulus onset and offset, but many occur during the intermediate period as well, especially in expert mice (Figure 2.11c, bottom). Temporal sequences of neuronal activity constitute an especially effective time coding scheme (Zhou et al., 2020); the fact that we see evidence of this type of temporal coding in primary sensory cortex supports the notion that timekeeping is a distributed process comprising multiple brain networks.

Movement does not account for learning-dependent changes in activity

Given that our behavioral paradigm is whisker-based, one might speculate that animals' whisking in response to stimulus presentation could modulate barrel cortex responses and could be responsible for the findings described thus far. To examine this possibility, we analyzed whisker motion from videos of the mice recorded throughout learning (Figure 2.13a). We quantified whisking motion as the mean difference between consecutive video frames. We identified bouts of whisking (including whisking during both trials and intertrial intervals) and calculated a whisk-triggered average of neuronal fluorescence (see 2.4   Methods). Comparisons of whisking-driven responses (Figure 2.13b, top) and stimulus-driven responses (bottom) of the same population of cells reveal that neurons respond strongly to stimuli, but not to whisking. In addition, the degree of whisking does not increase with learning: if anything, whisking decreases (Figure 2.13c).

Therefore, whisking cannot explain the aforementioned temporal information encoded by barrel cortex neurons. While the progression of time could be decoded from the activity of the cell population, it could *not* be decoded from whisker motion (Figure

2.13d). The vertical bars in Figure 2.13d show that decoders based on whisking consistently misclassify the time within the trial. Similarly, licking cannot account for the temporal decoding results, as we were able to decode time progression from population activity on "miss" trials during the stimulus (pre-reward) window, when there was no anticipatory licking (Figure 2.11a, upper-right panel). Thus, neither of these predominant movements can explain our results.
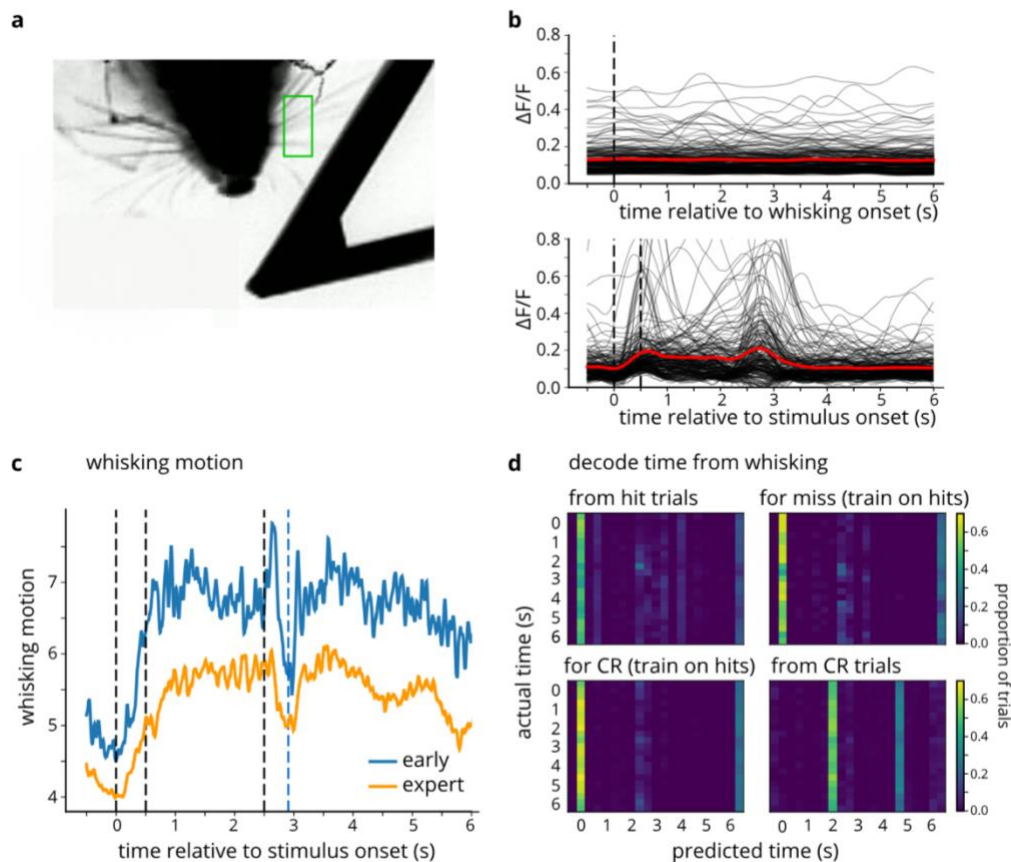
Figure 2.13: Whisking does not drive cells' responses, and is not responsible for learning-related increases in signal.

a) Frame from whisker video: whiskers are imaged from below. Green outline indicates ROI in which whisker motion was analyzed.

b) Top: whisk-triggered average for all cells (n = 3 mice); dashed vertical line indicates onset of whisk bout. Bottom: stimulus-driven responses for all cells (n = 210 cells). Red lines show average across cells. c) Average whisking motion on early and expert days (n = 3 mice).
d) Decoding time from whisking data.

Temporal surprise is an effective driver of neuronal activity

Finally, we investigated whether a mismatch of expectation and sensation can modify neuronal activity. Mismatch-related amplification of neuronal activity has been observed in primary visual cortex (Keller et al., 2012); in somatosensory cortex, on the other hand, cell responses were found to be dampened when sensory feedback did not match expectations (Ayaz et al., 2019). These previous studies investigated mismatch through perturbations of optic or tactile flow, by altering the velocity of the visual or tactile stimulus. Sudden shifts in velocity constitute a change of stimulus properties, which might alternatively explain the resulting variations in cortical activity (Muzzu and Saleem, 2021).

Given the ability of barrel cortex to represent within-trial time progression (as described above), perturbing the *timing* of the stimulus would be a useful means of disrupting an animal's expectations, without altering the intrinsic properties of the stimulus. Accordingly, we performed a *delayed-offset* experiment in expert mice, in which stimulus offset and corresponding reward were delayed by 1 second on 20% of the trials (Figure 2.14a, bottom). We compared trial-averaged fluorescence for delayed-offset vs normal trials and noticed that a large subset of cells (32%, 169/527) exhibited significantly

46

stronger offset responses on the delayed-offset trials compared to normal trials (see Figure 2.14a, top and middle for representative example cells). Note that we used a false discovery rate (FDR) correction to account for the large number of cells in these analyses. We quantified this difference across the neuronal population and discovered that a majority of cells show a greater response (39% increase on average) to delayed offset compared to normal stimulus offset (Figure 2.14b; Wilcoxon signed rank test, $p < 0.001$), indicating that surprising events yield stronger signals than expected events. Furthermore, the cell population response to delayed stimulus offset was on average even stronger (36% increase on average) than the response to stimulus *onset* (Figure 2.14c, $p < 0.001$), whereas for normal trials, the amplitude of the response to stimulus onset exceeded the response to stimulus offset (Figure 2.15a, $p < 0.001$). Thus, even in a primary sensory area, basic stimulus features are not necessarily the strongest driver of cortical activity.
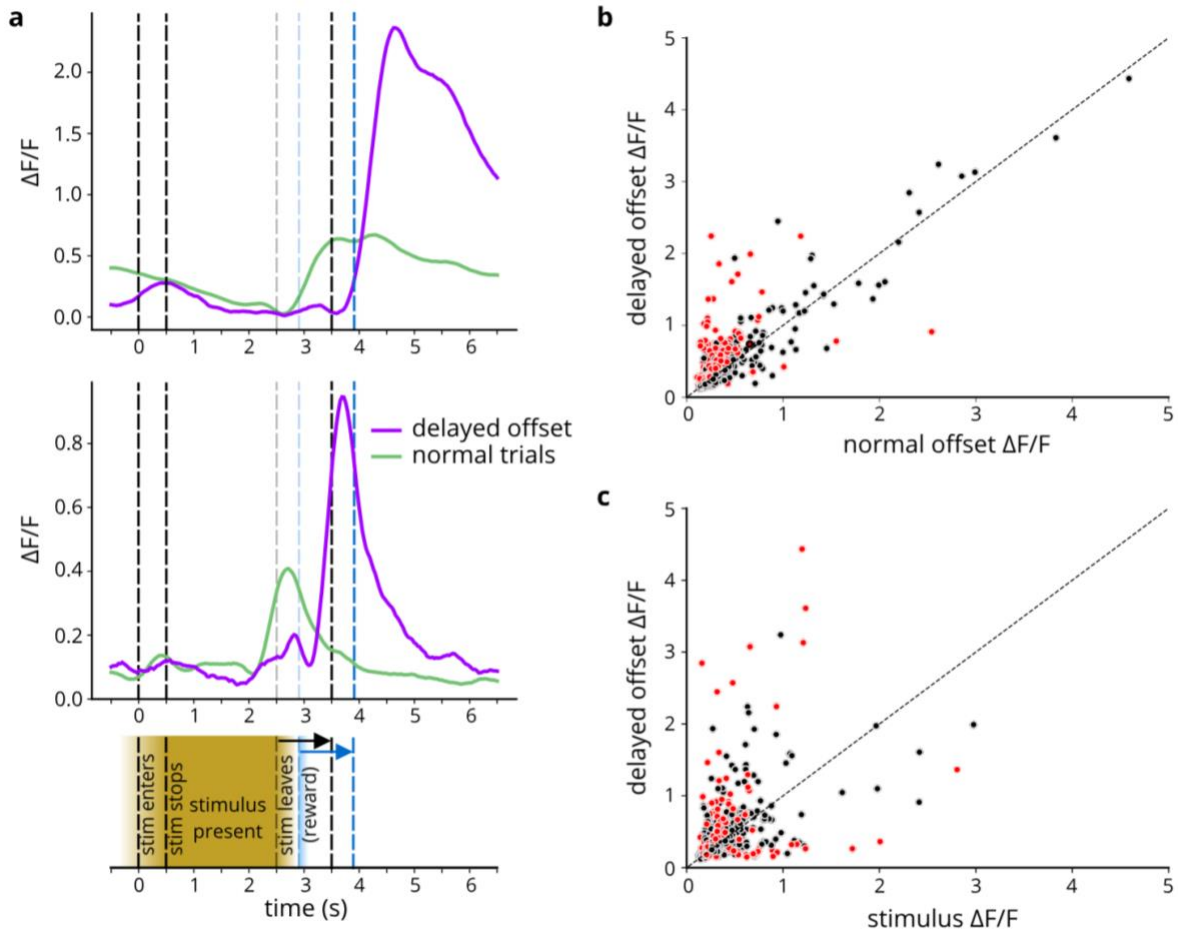
Figure 2.14: Delaying offset enhances size of late ("off") signal.
n = 3 mice (6 sessions)
a) bottom: schematic of behavior; top: 2 example cells' activity on normal trials (green; 80% of trials) and on delayed offset trials (purple; 20% of trials).
b) amplitude of late peak for normal vs delayed offset trials for all cells. Red dots indicate cells with significant difference (false detection rate corrected) in amplitude between normal and delayed offset trials. (normal vs delayed: p < 0.001).
c) amplitude of stimulus response vs offset response on delayed trials (p = 0.001).

We investigated whether the subset of neurons that preferentially responded to the unexpected stimulus offset in expert mice might have always responded to novelty by checking the responses of these cells on early training days, when all stimuli were novel. Interestingly, we discovered that these "surprise" cells did not respond to the stimulus

onset or offset in naïve mice (Figure 2.15b). On the surface, these observations may seem contradictory. However, the behavioral relevance of the stimulus offset on naïve and expert days differ markedly: on early-training days, stimuli may be novel, but are still meaningless to the mice, and are therefore not "surprising". From these results, we conclude that unexpected events are even more effective in eliciting layer 2/3 activity than sensory stimuli.
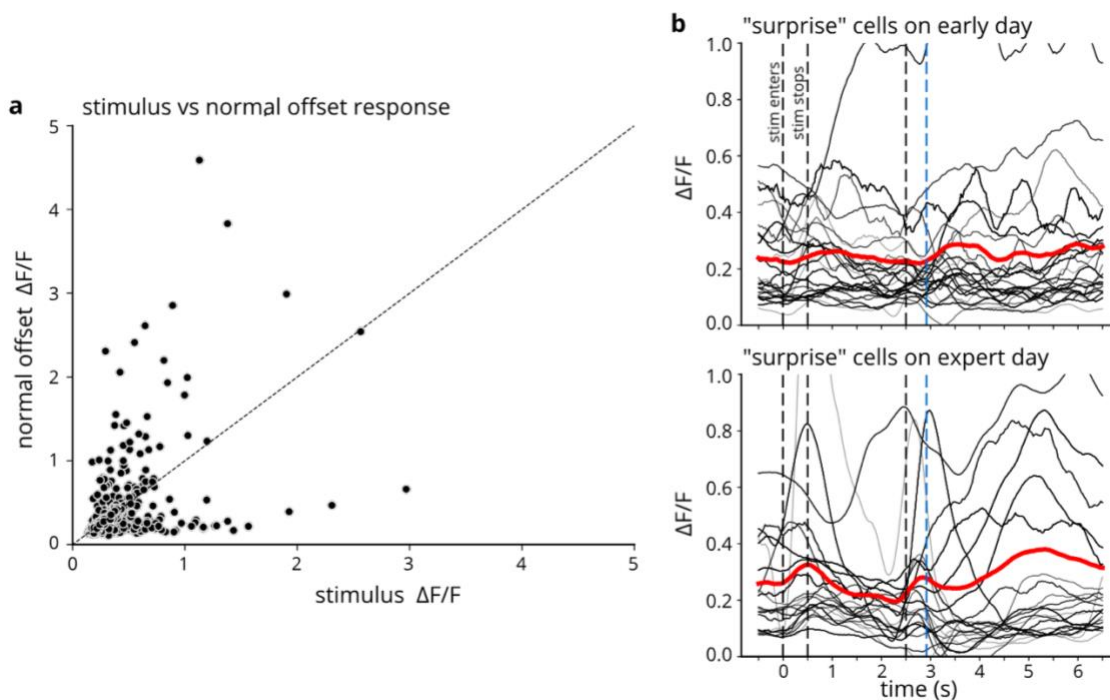
Figure 2.15: **Characterization of cells that respond to delayed offset.**
a) response amplitude to stimulus onset vs the response to stimulus offset for normal trials (p = 0.001).
b) longitudinally tracked "surprise" cells on an early training day (top) and on an expert day (bottom). Red trace is the mean across cells.

2.3 Discussion

In recent years, a growing body of research has begun to reveal the nuances of processing in primary sensory cortices. Contrary to the canonical view of primary sensory cortex as an early stage in a sensory processing hierarchy, we now have reasons to believe that primary cortex may have a role in multisensory integration (Budinger et al., 2006; Ghazanfar and Schroeder, 2006; Maruyama and Komai, 2018; Zhang et al., 2020) and in the processing of non-sensory information such as choice, reward, and reward expectation (Lacefield et al., 2019; Pleger et al., 2009; Rodgers et al., 2021; Shuler and Bear, 2006). In addition to this evidence for complexity and higher-level processing in primary cortical areas, parallel evidence contradicts the straightforward notion that primary cortical cells are highly responsive to stimuli of the corresponding modality: for instance, while deep layer cells in barrel cortex respond vigorously to whisker stimulation, even spatially and temporally complex whisker stimuli elicit only low levels of activity in superficial layer cells (Estebanez et al., 2012; Ramirez et al., 2014), though these cells may respond more to *behaviorally relevant* whisker contacts (Rodgers et al., 2021). The current study contributes to this alternative view of primary sensory cortex as akin to association cortex, and unveils additional facets of cortical processing.

Here, we demonstrate that learning recruits previously unresponsive cells in barrel cortex, creating a neuronal population that better represents tactile stimuli that have been paired with reward. Previous studies have identified structural and functional plasticity in barrel cortex after sensory deprivation (Chau et al., 2014; Margolis et al., 2012) as well spine

plasticity after learning (Kuhlman et al., 2014), so one might have anticipated learning to alter neural representations as well. However, there remains disagreement in the field on this issue: some studies have indeed found learning-related functional changes (Chen et al., 2015), while other studies observed more stability across learning. Kim et al. (2020) observed a high variability and turnover in the responsiveness and selectivity of barrel cortical cells while mice learned an object-angle discrimination task, but those authors and others (Kim et al., 2020; Peron et al., 2015) reported that the proportion of responsive cells remained unchanged across learning. Interestingly, Makino & Komiyama (2015) find that the number of layer 2/3 responsive cells *decreases* with learning (Makino and Komiyama, 2015), though that study is in a different modality, and in an aversive conditioning task, which may explain the discrepancy. In contrast with these latter studies (but in agreement with Chen et al., 2015), we find that the ratio of responsive neurons markedly increases with learning: this recruitment of previously unresponsive cells into a newly responsive neuronal population explains our observation of increased overall touch responsiveness of barrel cortex. As for individual longitudinally tracked cells, we did observe some degree of bidirectional turnover, including both cells that lost responsiveness as well as those that became more responsive. However, training appeared to rebalance these dynamics such that, while only a minority of cells that were originally active lost their responsiveness, most unresponsive neurons gained stimulus responses.

In stark contrast, repeated exposure to a stimulus without reward-pairing actually *reduced* barrel cortex responsiveness to the repeatedly presented stimulus. The latter phenomenon likely arises due to habituation: whereas mice learn the behavioral relevance

51

of a stimulus followed by reward, they can equally learn the relative unimportance of a stimulus that is presented without positive or negative consequences. Similar findings have been reported in the mouse primary visual cortex (Henschke et al., 2020; Keller et al., 2017)—repeated presentation of visual gratings led to adaptation of neural responses over the course of behavioral training.

One plausible mechanism for this retuning of sensory representation is contextual information conveyed to barrel cortex via long-range top-down inputs from higher-order cortical areas. Orbitofrontal cortex (OFC) has recently been shown to play an important feedback role, but is unlikely to be involved in the effects described here: input to S1 from OFC has been shown to be activated (and indeed required) in complex tasks such as reversal learning, but OFC only became active during the complex (rule switch) aspects of the task, not during initial learning (Banerjee et al., 2020). However, other top-down inputs, such as those from the prefrontal cortex (PFC), could contribute in our task and others (Fritz et al., 2010; Hamm et al., 2021; Rodgers and DeWeese, 2014). For instance, some high-level information about task context has been identified in both primary auditory cortex (A1) and PFC (Rodgers and DeWeese, 2014) during auditory behaviors— presumably, this information was conveyed to the primary sensory region from PFC. In addition, prefrontal inputs to primary visual cortex have been shown to enhance mismatch signals (Hamm et al., 2021). Even more explicit is evidence for top-down modulation of primary sensory cortical activity by retrosplenial cortex during learning (Makino and Komiyama, 2015).

Alternatively or additionally, the effects we observed may have a mechanism of neuromodulatory origin, whereby in the presence of reward, neuromodulatory inputs to barrel cortex create an environment for local plasticity to occur and unveil neuronal responses to the stimuli. While multiple neuromodulatory systems may be implicated, several are particularly compelling. Norepinephrine, for instance, is intimately involved in task engagement and attention, and could play a role in the enhancement of sensory responses to behaviorally relevant stimuli. In fact, there exists evidence for the ability of noradrenergic neuromodulation to induce plasticity in primary auditory cortex and strengthen A1 responses to auditory stimuli (Martins and Froemke, 2015) and to increase neuronal excitability in S1 (Labarrera et al., 2018). Acetylcholine, too, has been shown to augment neuronal responses in studies of V1 (Goard and Dan, 2009; Pinto et al., 2013). The fact that reward uncoupled with stimulus presentation is sufficient to potentiate cell responses in V1 (Henschke et al., 2020) supports the neuromodulation model, as the slow time course of neuromodulation may allow for the potentiation of signals even when those signals are not temporally locked to reward. Future studies should investigate this possibility and identify which neuromodulators might be responsible for retuning cortical representations according to reinforcement.

Our results also illustrate how barrel cortex activity becomes more patterned as mice learn the behavioral task, as confirmed by the improved ability of a classifier to decode the presence or absence of the stimulus as well as the nature of animals' responses. Furthermore, learning enabled decoding of the trial type from more time bins and, crucially, *earlier in the trial*, allowing cells to more quickly access information about the

53

trial. This phenomenon may underlie enhanced reaction times as behaviors become more ingrained, and may be explained by a temporal difference model of reinforcement learning, whereby the reward value is attributed to progressively earlier stimuli (or perhaps to earlier moments of time within the trial) over the course of learning.

Barrel cortex encodes time progression

We found that time could be decoded with remarkably high precision from population activity in well-trained mice but not in naïve animals, meaning that cells in barrel cortex could track the progression of time. Previously, in barrel cortex, temporally precise response patterns have been shown to carry sensory information (Arabzadeh et al., 2006), though this finding was obtained in anesthetized rats, making the relevance to behavior and learning somewhat difficult to interpret. However, a related finding in the somatosensory cortex of awake primates discovered so-called "memory cells", which displayed specific patterns of activity during the delay period of a tactile working memory task (Bodner et al., 2005). However, our ability to decode time likely stems from time cell-like responses across the neuronal population rather than from temporally precise activity patterns of individual cells like those observed by Arabzadeh et al. To our knowledge, our results present the first direct evidence of activity in barrel cortex encoding the progression of time on a behavioral timescale, or of temporal encoding in any rodent primary sensory cortical area.

One could imagine that the temporal information in S1 cells arises due to patterns of whisking: stereotyped whisking patterns would lead to temporally precise series of whisker contacts; consequently, even though the stimulus is not moving, and seemingly no new events occur, a series of stimulus responses could emerge. The time course of such signals could explain the source of the timing information as well as our inability to decode time on stimulus-absent trials. However, we were not able to decode time from whisking (Figure 2.13c), ruling out whisker movement as the relevant source of temporal coding in S1.

Cells respond to mismatch of temporal expectations

Our experiments uncovered mismatch signals in barrel cortex: in expert mice, delivering the stimulus offset at an unexpected time yielded the highest levels of cellular activity. Unexpected events have previously been shown to evoke strong responses in primary sensory cortex (*e.g.*, oddball stimuli eliciting responses in "deviance detector" cells of V1 (Hamm et al., 2021)). Those kinds of experiments, however, typically use stimuli that have different sensory characteristics from the "non-deviant" stimuli. As a result, the increased responses to oddball stimuli can be explained by synaptic or circuit adaptation to the common stimulus, and a lack of adaptation to the oddball.

Other studies (Ayaz et al., 2019; Keller et al., 2012) have reported altered responses in V1 and S1 to perturbations to optic and tactile flow, respectively, in which sensory input did not match the expected outcome of self-generated motion. Yet, an alternative

explanation for these results has been offered (Muzzu and Saleem, 2021), invoking cortical cells' preferences for different stimulus velocities, rather than mismatch between expectations and sensory feedback. An important advantage of our experimental design is that we do not change velocity or any other stimulus property to investigate mismatch, but rather the timing of the stimulus. Consequently, the surprising events in our experiment have the same sensory features as expected events, yet they violate the animal's internally constructed model of the world—the temporal structure of the task.

In fact, the "surprise" signals we observe constitute further evidence for the existence and importance of timing information in S1: the enhancement of cells' responses to a disturbance of an event's timing implies that they have "knowledge" of its time course, and access to timing information is necessary for sensory cortical cells to exhibit this response to temporal mismatch. The presence of a surprise response in primary sensory cortex may in turn provide the animal with a means for rapid behavioral outcomes to salient, unexpected events, perhaps bypassing slower processing in more traditional association cortical areas.

In conclusion, we have demonstrated that primary somatosensory cortex behaves in a more nuanced way than traditionally thought: rather than directly responding to sensory stimuli, sensory cortical representations undergo large bidirectional changes over the course of learning, either expanding to better represent stimuli paired with reward, or habituating and contracting following repeated exposure without reward pairing. Indeed, these effects extend beyond mere stimulus responses, and encompass temporal encoding

as well as modelling expected events and their time course, implying that timekeeping may be widespread and distributed throughout the brain.

2.4   Methods

Subjects

All experiments were approved by the Columbia University Institutional Animal Care and Use Committee. 18 C57BL/6J mice were used in the experiments described here. Mice were housed in groups of 2-5, unless fighting or barbering was observed (in which case they were singly housed). They were provided with a running wheel for enrichment as well as *ad libitum* food. During behavioral training, mice were water-restricted and maintained at 80% of their original weight. If daily weighing determined that mice were too light, they were given 5 minutes of free access to water in their home cage.

Virus injection and cranial window implant surgeries

CaMKII-GCaMP6f virus (AAV5.CamKII.GCaMP6f.WPRE.SV40, nominal titer 2.3 x 10^{13} gc/mL) was injected into left barrel cortex of the mice. The injection site was targeted at 1.5mm posterior and 3.5 mm lateral relative to bregma. To accomplish this surgery, mice were anesthetized with isoflurane (3% for induction; 1-2% for maintenance), and subcutaneous analgesics buprenorphine, carprofen, and bupivacaine were administered.

Eye ointment was applied, and mice were fixed into the stereotax with earbars. Scalp was shaved and cleaned, and then cut to expose the skull. After locating the injection site, a small area of skull was thinned with a dental drill, until a glass injection pipette could be inserted through cracks in the thinned bone. Virus was injected at depths of 300 μm and 150 μm, with three 50-nL injections at each depth (injections were spaced 1 minute apart, with three minutes between depths and before removing the pipette). The pipette was then withdrawn, the thinned area of skull was covered with superglue, and the scalp was sutured closed.

Mice were given an additional carprofen dose 24 hours after surgery, and monitored for 5 days. We allowed a three-month interval for the virus to express, and then performed a cranial window implantation surgery.

For this second surgery, intramuscular dexamethasone was administered 3 hours prior to surgery. Isoflurane anesthesia was carried out as before, and buprenorphine was administered subcutaneously. Remaining surgical preparation was done as before, but this time a circle of scalp was removed entirely. The skull was cleaned, and a 4-mm diameter craniotomy was made over the barrel cortex. A glass cover slip was placed over the craniotomy and sealed with superglue. Then a metal headplate was affixed to the skull using dental cement. Subcutaneous buprenorphine was given every 12 hours for two days after surgery. Mice were allowed two weeks to recover and for blood to clear from the cranial window.

Intrinsic signal imaging and 2-photon calcium imaging

Barrel cortex was located using intrinsic signal imaging in anesthetized mice, during which individual whiskers were stimulated with a piezoelectric device at 5Hz (10 pulses, with a 10-second gap between pulse trains). Under red light illumination, and through a 5x objective, we recorded reflectance at the brain surface with a Rolera-MGi Plus digital camera: changes in blood flow to the barrel corresponding to the deflected whisker yielded a visible change in reflectance.

Two-photon calcium imaging, through a 16x /0.8NA water immersion Nikon objective and with the laser set to 940 nm, was conducted on a Sutter Moveable Objective Microscope; images (512 x 512 pixels) were acquired using ScanImage software with a 30Hz acquisition frame rate.

Cells in L2/3 of barrel cortex were imaged each day during behavioral training, and the field of view was kept constant between days. The center of the barrel field (C or D row) was targeted for imaging.

Whisker-based object-detection behavioral paradigm

Our behavioral paradigm was a Pavlovian object-detection task. First, mice were hand-habituated for several days. Second, a "lick-training" phase occurred: mice were water-restricted and head-fixed into the behavioral apparatus, and water was delivered through a lick-port. Initially, water was delivered manually; once mice began to

spontaneously lick at the lick-port, water was delivered contingent upon the animals' licking. Licks were measured using an infrared lick-detector: whenever the tongue touched the lick-port, it covered a fiber-optic cable leading to the infrared detector, and the event was registered as a lick.

Finally, the detection task training began. In this paradigm (Figure 2.1a), two objects were affixed to opposite arms of a plus-shaped wheel. The other two arms were left empty. A motor rotated the wheel, and brought one of the arms toward the animal's face on each trial. The object "enters" the whisker field at time 0 (Figure 2.1a, bottom), and then stops 500ms later. The wheel remained in this position for two seconds, and then rotated 45 degrees, such that the space between two arms was in front of the face. On each trial, the wheel randomly rotated either clockwise or counter-clockwise; the arm that was ultimately presented was also selected at random. If the arm presented to the mouse had an object attached to it, that trial was a "stimulus-present" trial, and was followed by water reward. The onset of water delivery began 407ms after the wheel began to rotate away (10ms after the wheel completed the 45 degree rotation). On unrewarded, "stimulus-absent" trials, one of the empty arms was presented, and no water was given. The aforementioned infrared lick-detector was used to measure licks: anticipatory licks were counted during the two seconds leading up to the reward, or lack thereof.

For the *delayed-offset experiment*, stimulus onset did not change, but for a random 20% of trials, the offset was delayed by 1 second (*i.e.*, for those trials, the stimulus (or empty arm) was present for 3 seconds rather than the usual 2 seconds). On rewarded trials, the reward timing did not change relative to the offset of the stimulus.

To ensure that any effects we observed were due to stimulus-reward conditioning rather than repeated exposure to stimuli, a control group of mice were trained without rewards. These mice had *ad libitum* access to water, and during behavioral training, no water was delivered through the lick-port. The rest of the task parameters remained the same, and these mice were trained, on average, for the same length of time as the mice in the conditioning cohort.

Histological confirmation of imaging location

After training and imaging experiments were complete, the two-photon laser was used to create a small lesion at the center of the imaging field of view (at a depth of 100-150um). Mice were then perfused and their brains harvested. 50um tangential sections were made, and stained with streptavidin-Alexa 647 in order to visualize the barrels.

Data processing and analysis

Motion correction and cell identification:

Suite2p was used for motion correction, to identify ROIs, and to extract their signals (Pachitariu et al., 2016). These automatically detected ROIs were them manually curated to remove likely false positives.

Longitudinal tracking:

Motion-corrected time averages were aligned across any two days by warping one image (time average of Day 2) to match the other (time average of Day 1) with an affine transformation (using the Python OpenCV library). The same transformation was then applied to each ROI mask of Day 2, such that the ROI was warped and shifted to the appropriate location. All ROI locations from the two days were then compared and checked for degree of overlap. Overlapping ROIs were given the same label, which identified them as corresponding to the same cell across the two days. In cases where an ROI from one day overlapped with more than one ROI from the other day, ROI pairs with the greatest degree of overlap were given the same labels. Once a cell has been found in two imaging sessions and given a label, it maintains that identity when tracked across additional sessions: thus, a cell can be longitudinally tracked across numerous sessions.

Quantifying cell responses:

Signals were smoothed with a Savitzky–Golay filter. ΔF/F at each timepoint was calculated using a baseline determined by the 8th percentile of a 50 second rolling window centered at that timepoint.

Transient detection:

For each cell, a threshold was set at 2 standard deviations above the median $\Delta F/F$ for that cell. A transient event onset was marked whenever the $\Delta F/F$ first crossed the threshold (and the transient event was considered to end when the $\Delta F/F$ fell below the threshold). (See Figure 2.2.)

Classifying cell response types:

Cells were classified as "on", "off", "on-off", "reward", or "none" cells based on the timing of their responses. Once transients were detected for each cell, a histogram was calculated indicating the probability of a transient event occurring at each time within a trial; this probability curve was then compared to that of shuffled data. Shuffled data was generated by doing 5000 iterations of a circular shuffle, whereby the data in each trial was shifted by a random number of frames; thus, any temporal dynamics were preserved, while the relationship of the data to the trial timecourse was disturbed. A Z-test compared the transient probabilities during relevant 1 second intervals (around the stimulus onset and offset, and following reward) to the shuffled data on the same intervals. This test revealed time intervals during which the probability of a transient occurring was greater than chance. If a cell's transient probability was significantly greater than chance only during the stimulus onset bin, that cell was classified as an "on" cell; cells with significant transient probability only at stimulus offset were "off" cells; "on-off" cells had significant transient

probability at both intervals, and "reward" cells had significant transient probability after reward. A false discovery rate (FDR) correction accounted for the large number of cells.

Decoding analyses:

We trained linear support vector machine (SVM) classifiers to decode stimulus (*i.e.*, classify trial type as "stimulus-present" or "stimulus-absent"), and to decode choice (*i.e.*, classify the animals' response as "lick" or "no lick"), using 4-fold cross-validation (training on 3/4 of the data and testing on the remaining 1/4). In both cases, we matched the number of trials between the two classes. We computed the separability of the classes by projecting (via dot product) neuronal activity of trials of each class onto the coding direction (as determined by the classifier weights), and quantifying the earth mover's distance between the distributions. Overlap between the distributions corresponds to decoding errors, and the classes are judged to be more separable the further apart the distributions are.

To decode time, we split trials into 10-frame-long time bins (1/3 second duration), and trained SVM decoders to distinguish between time bins. A confusion matrix (see Figure 2.11a) with true time bin identity on the y-axis and predicted time bins on the x-axis illustrated the ability to decode time, by showing the proportion of trials for which a given time bin y was predicted to be time bin x. The matrix was normalized by the true time bin classes, so the values of each row sum to 1.

Identifying temporal tuning:

The degree of temporal information in each cell's activity was quantified using the Skaggs spatial information metric (Skaggs et al., 1993). This metric is commonly used to quantify spatial tuning, but has been applied to the temporal dimension as well (Shimbo et al., 2021). We defined cells as time-tuned if their temporal information exceeded the 95[th] percentile of the temporal information in shuffled data. Among these time-tuned cells, we then identified the specific temporal "regions" to which they were tuned, where each cell's "tuning curve" (probability of transients across time) was above 95% of shuffled curves.

Whisking analysis

Throughout conditioning, whisking videos were acquired at 125 Hz using a Sony PS3eye camera. Using the Python OpenCV library, we extracted pixel values within a hand-drawn ROI near the animals' face, on the side of the face contacted by the stimuli. The mean difference between consecutive video frames was used as a measure of whisking motion. Whisking bouts were identified in a similar manner to calcium transient events, by tracking where the whisking motion crossed a threshold. We defined a whisking bout as lasting >0.5 seconds, so a new bout could not begin until that time interval elapsed. Whisk-triggered averages were computed by aligning fluorescence data with the onset of whisking bouts.

Chapter 3 : Feature discrimination—a complex behavior

3.1   Introduction

In Chapter 2, I discussed the effects on cortical activity of learning a very simple object detection task. The simplicity of the behavioral paradigm certainly had benefits, such as rapidity of acquisition and minimization of potential confounds. However, the object detection task was very possibly not cortex dependent: other object detection paradigms have been found to not require cortex for learning or execution of the behavior (Hong et al., 2018). In that case, probing cortical activity in a more complex or ethological task may be informative.

One option would be to use an ethological task known to depend on cortex: the "gap cross task" is a freely-moving behavioral paradigm which requires animals to use their whiskers to determine the width of a gap between two platforms, and decide whether they can safely jump from one side to the other (Hutson and Masterton, 1986). Although similar to an object-detection task (since animals have to detect the presence of the platform on the other side of the gap in order to decide that it is safe to jump across), the gap-cross task is more complex than a simple whisker-based object-detection task, because mice have to integrate information such as location and head position in order to determine whether they can cross the gap. I did work on developing a mouse version of this task (Figure 3.1), which was originally performed in rats, but in this chapter, I will focus on a different paradigm—a complex object discrimination task, described below.

66
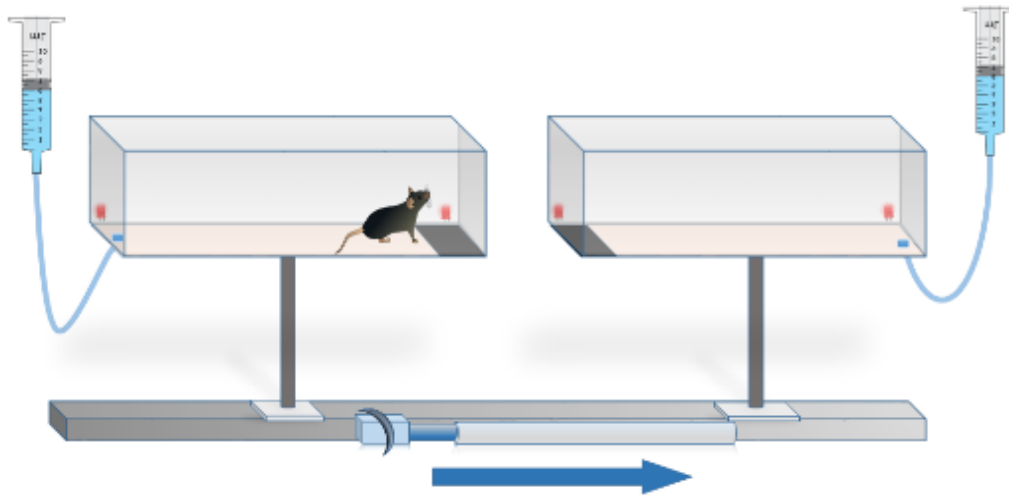
Figure 3.1: Gap cross task.
Mouse jumps between two elevated platforms, with variable gap width.

Most object discrimination tasks ask animals to use a single feature type to distinguish between objects. The more complex of these tasks may use gradations of the feature—*e.g.*, the roughness of a texture (Park et al., 2022) or the orientation of a visual stimulus (Lyamzin et al., 2021), etc. However, fewer studies have examined object discrimination where each object is defined by multiple features. In a bi-conditional or "exclusive-or" (XOR) task, there is a non-linear mapping of stimulus to animals' response, making the task much more complex: the presence or identity of a given feature is insufficient for the animal to make a choice. Figure 3.2 shows an example of a visual XOR task, one performed by monkeys (Anderson et al., 2006).

Right
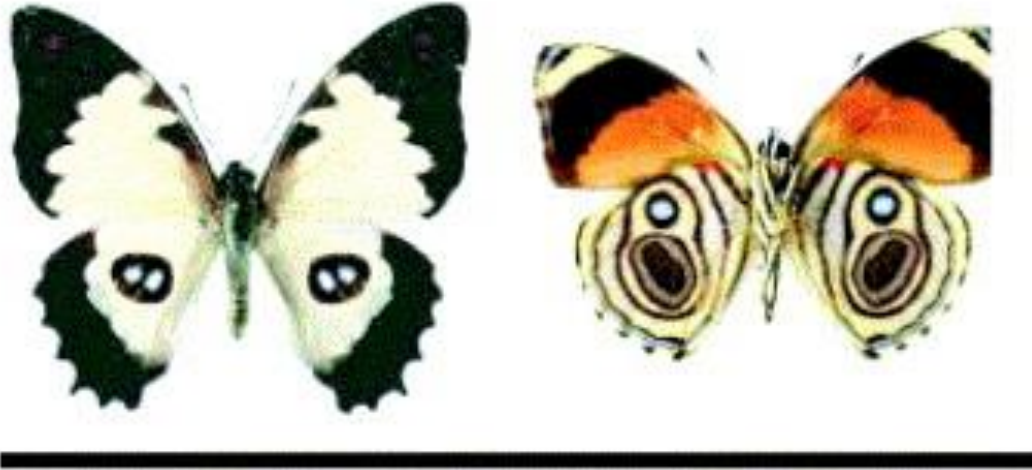
Left

Figure 3.2 Example of a bi-conditional (XOR) task
Adapted from Anderson et al. (2006). In these stimuli presented to monkeys, each object is defined by the overall butterfly shape/color/pattern (large white butterfly vs smaller orange butterfly), as well as the specific pattern on the lower part of the wings (two small horizontally oriented white dots per wing, or one diagonal dark patch). The animal's response (in this case, a left or right button press by the monkey) is guided by the conjunction of both features.

While generally thought to be difficult for animals to acquire these tasks, XOR tasks have been occasionally been successfully carried out (Anderson et al., 2006; Ramirez and

Colwill, 2012). For instance, mice were able to learn an XOR task when each condition was defined by an auditory cue and an environmental context; however, when each condition was defined by multiple discrete cues rather than a context (in this case, an auditory and a visual cue), mice were unable to acquire the task (Ramirez and Colwill, 2012).

I chose to administer a tactile XOR task to mice: while I was concerned about the level of difficulty, the complexity of the task made it a likely candidate for a cortically-dependent behavior.

3.2   Results

I developed a complex object discrimination task (Figure 3.3), in which head-fixed mice learned to use their whiskers to discriminate four objects, on the basis of conjunctions of features—shape (concave vs. convex) and texture (smooth vs. rough). Of these four objects, two were rewarded, and two were unrewarded; the two defining features of a rewarded object could both also be found on unrewarded objects (but not together on the same unrewarded object). Thus, no one feature alone could provide sufficient information about whether reward would follow a given object, and mice had to base their decisions on both features.

Curvature

convex   **concave**

convex rough: concave rough:
rewarded   unrewarded

concave rough: unrewarded...

convex smooth: concave smooth:
unrewarded   rewarded
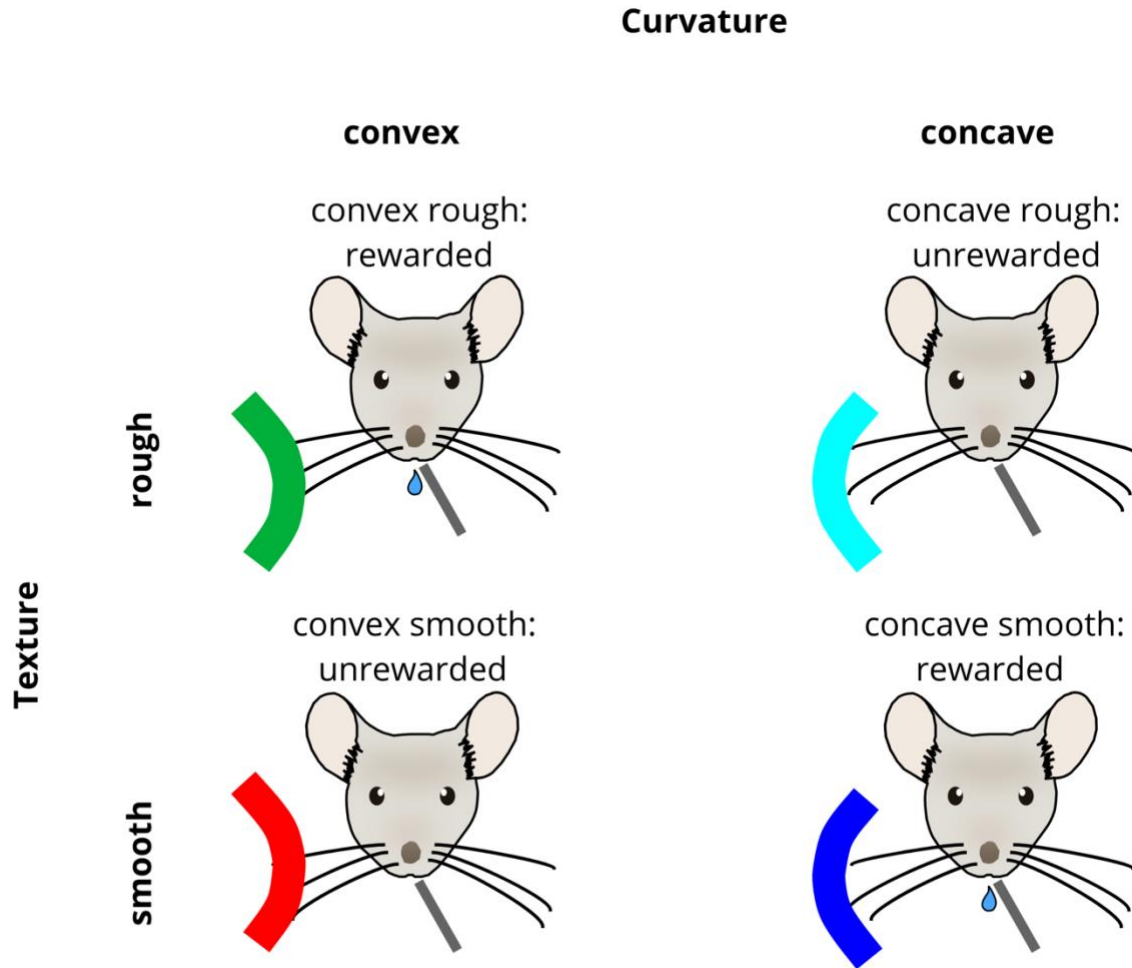
Figure 3.3 Schematic of behavioral paradigm.
Curved objects are presented to the mouse. Convex rough (green) and concave smooth (blue) objects are rewarded. Concave rough (cyan) and convex smooth (red) objects are unrewarded.

Preliminary results suggested that mice were able to learn this task quickly: within two weeks of training, 10 out of 19 mice began to consistently show anticipatory licking preferentially to the rewarded objects. Figure 3.4 shows the performance of an example mouse: the solid blue and green lines correspond to rewarded trials, and the dotted red and cyan lines correspond to unrewarded trials. Early in training, all four lines hover near each

other, indicating similar rates of anticipatory licking for all trial types. Then, at session 15, the lines rapidly diverge, implying that the animal perhaps had a sudden realization concerning the task rules.
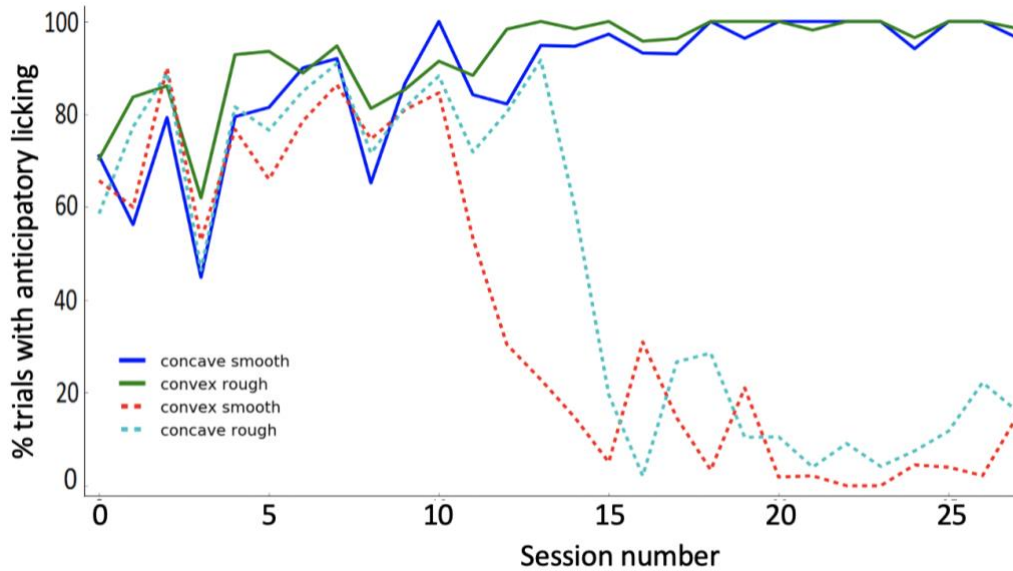
Figure 3.4: Example mouse performance, as measured by anticipatory licking
Solid lines (blue and green) correspond to rewarded trials; dotted lines (red and cyan) correspond to unrewarded trials.

However, it soon became apparent that this mouse and many others were solving the task without using their whiskers, presumably by relying on other senses: after having their whiskers trimmed, their performance was frequently unimpaired (see example mouse in Figure 3.5).
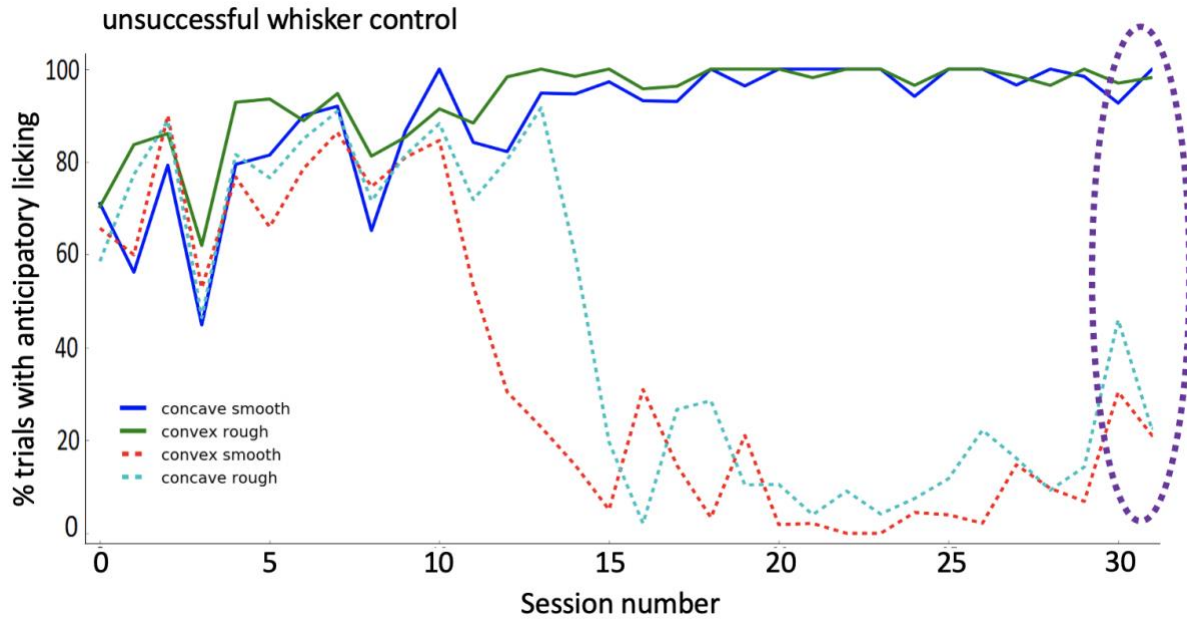
Figure 3.5: Example behavior of a mouse relying on modalities other than whisking.
Solid lines (blue and green) correspond to rewarded trials; dotted lines (red and cyan) correspond to unrewarded trials. The two last sessions, encircled by the purple dotted ellipse, are the whisker-trimmed sessions. Note that the solid and dotted lines remain separated on the whisker-trimmed sessions.

In an attempt to dissuade animals from using other modalities, I made the following changes to the experimental design. First, I introduced an initial phase of training in which the behavioral apparatus was illuminated. I speculated that the ability to see the objects while being familiarized with the task would help mice understand the actual task structure and prevent future "cheating". In fact, there is a precedent for training whisker behaviors in the light first, to acclimate animals to the training environment (Morita et al., 2011; Zuo et al., 2011). Of course, I did not want mice to rely on vision, so once mice were consistently

72

able to discriminate the objects, the illuminating lamp was turned off during the presentation of the shapes, and turned on as a blinding light during the inter-trial interval to prevent any dark adaptation (in case there was any light leak that might enable mice to rely on their vision). To control for olfactory cues, I switched from using plastic objects (which were used in pilot experiments) to aluminum ones which I wiped daily with ethanol: metal should absorb fewer odors, and should be easier to clean.

Despite all these efforts, animals' seeming insistence on circumventing these measures was undeterred. Of 17 animals that learned the task, only 3 mice displayed a drop in performance following whisker trim, even after all these measures; only a single animal's performance decreased to *chance* levels after trimming whiskers (Figure 3.6). Consequently, I conducted the following control experiments, testing possible methods by which the remaining animals might be "cheating".

Possible sources of animals' cheating abilities

Position of objects rather than sensory features (also, magnetosensation)

From the outset of these experiments, the experimental paradigm was designed to maximize randomization. The direction of rotation of the wheel was randomly selected on each trial, as was the final position. Thus, from a given position, the wheel could rotate to any of the 4 positions that corresponded to the four trial types. And since the rotation direction was randomized, both a short and a long rotation could lead to the same object

being presented. These characteristics of the experimental design should have made it difficult for animals to identify a given object based on its spatial and temporal relationship to the previous one. Nevertheless, some aspects of the experimental design were not randomized: each session began with the wheel in the same starting position, and the objects always maintained the same spatial relationship to each other, as they were in fixed locations on the wheel.

I also considered that mice might be able to determine the position of the wheel via magnetosensation: after all, some evidence does suggest that mice and other mammals are capable of magnetosensation (Mather and Baker, 1981; Nemec et al., 2001). In early iterations of the experiment, I had originally placed magnets on two of the wheel's arms: in combination with a sensor, the magnets were supposed to help the wheel stop at the correct position. I wondered whether mice might be sensing the magnets, and using them to determine the position of the wheel.

So, might animals be able to correctly identify trial type from the *position* of the wheel, rather than from the sensory *features* of the objects themselves?

To test this question, we removed the objects entirely, or re-ordered the objects to destroy any learned spatial relationship between them. With no objects present, animals were completely unable to report whether a trial was going to be rewarded or unrewarded. And with objects repositioned, animals continued to correctly lick in response to the convex rough and concave smooth objects and suppress licking to the convex smooth and concave rough ones, even though the objects were no longer in their original positions.

Thus, mice did in fact learn something about the objects themselves, even if it was not their tactile features.

Audition

I suspected mice could be relying on the ambient sounds altered by the shape of the objects next to the ear—similar to how a seashell held against the ear yields an auditory stimulus. To test whether mice were using auditory cues to identify the objects, we first tried physically blocking the ears by plugging them with latex. However, I soon realized that this control experiment was likely flawed: blocking animals' ears with latex certainly impaired performance of whiskerless mice, but probably only because the mice were too distracted by their new earplugs. We quickly abandoned this disruptive and distracting approach. Less invasive controls—such as playing white noise out of a speaker to mask any auditory cues—were unsuccessful, with whisker-trimmed mice still able to perform the task.

Olfaction

Despite daily cleaning of the metal shapes to remove lingering odors, it was possible that mice still used olfactory cues. Indeed, when I soaked the objects in ethanol overnight, the performance on the task dropped dramatically. Further, whiskerless mice were

completely unable to do the task after I added a fan to direct airflow (and potential odors) away from the mouse.

Unfortunately, following the implementation of this new control (the fan), trained mice were unable to perform (or relearn) the task even with a full set of whiskers. Nor were new cohorts of mice ever able to learn the task.
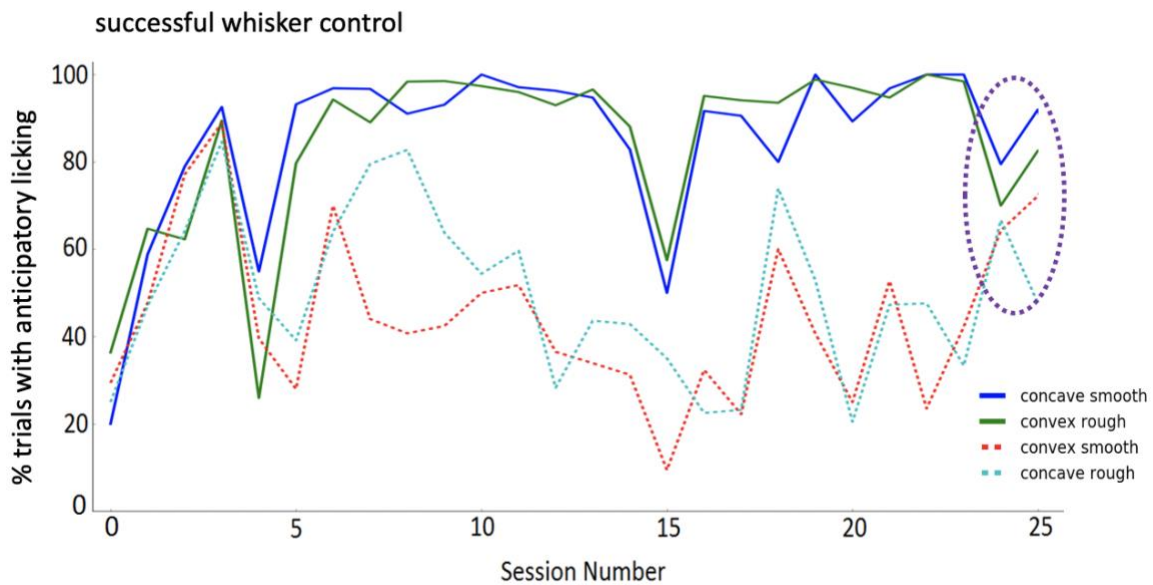

Figure 3.6: Example behavior of a mouse that likely uses whiskers to perform the task.
As in Figure 3.1, solid lines correspond to rewarded trials and dotted lines correspond to unrewarded trials. The last two sessions, circled, are the whisker-trimmed sessions. Note the convergence of the dotted and solid lines toward each other.

3.3   Discussion

A total of 17 out of 27 mice seemed to learn the object discrimination task; yet, I could only be confident that a single mouse ever learned it as intended, entirely relying on whiskers (Figure 3.6). I was forced to conclude that this complex behavioral task was simply too difficult for most mice to learn using whiskers alone.

The fact that a single mouse was able to correctly acquire the task demonstrates that mice may in principle be able to learn a tactile XOR task. In addition, two other mice did exhibit a more modest drop in performance after having their whiskers trimmed, suggesting that they did partially use tactile information, but supplemented with other modalities. However, the experimental design as currently implemented might be too complex for mice to reliably learn for experiments requiring large cohorts. We have not ruled out the possibility that the experimental design could be optimized further to improve learning. Several possible changes could be implemented to make it easier for mice to use their whiskers to solve this feature discrimination task. Most notably, the texture used was rather subtle, and could be replaced by a more noticeable texture. Other studies in our lab and elsewhere successfully trained mice to discriminate textures (Chéreau et al., 2020; Park et al., 2022), but opted for different texture design. One option is to use sandpaper and to mitigate any odors the sandpaper absorbs by using a fan (as described above); another is to make the grooves more pronounced in depth and placed further apart. Future versions of the task can implement these changes, as well as being vigilant about olfactory cues.

3.4   Methods

Subjects

Same as in Chapter 2. All experiments were approved by the Columbia University Institutional Animal Care and Use Committee. 31 mice were used for these experiments.

Headplate surgery

Mice were anesthetized with isoflurane (3% for induction; 1-2% for maintenance), and subcutaneous analgesics buprenorphine, carprofen, and bupivacaine were administered. Eye ointment was applied, and mice were fixed into the stereotax with earbars. Scalp was shaved and cleaned, and a circle of scalp was removed to expose the skull. The skull was cleaned and scored with a scalpel, and metal headplate was affixed to the skull using dental cement. Mice were given an additional dose of carprofen 24 hours later.

Object discrimination behavioral paradigm

We trained mice on a Pavlovian object discrimination task. Pre-training was as described in Chapter 2, and consisted of habituation and lick-training. Licks were measured using a capacitance lick-detector.

For the object discrimination task, one of four objects was affixed to each of four arms of a plus-shaped wheel. The rotation and timecourse of the task were identical to those described in Chapter 2.

The four objects were curved aluminum shapes, either concave or convex, and either rough or smooth. 100um deep grooves with 250um spacing that were milled into the surface of the objects provided the rough texture. The curves had a 36mm radius. The convex rough object and the concave smooth object were both rewarded and were followed by a water reward; the convex smooth and concave rough shapes were unrewarded.

Anticipatory licks were counted during the two seconds that the object was stationary in front the animal's face.

Whisker trimming control

In order to test whether mice were using their whiskers to solve the task, or relying on some other modality, a whisker-trim control was performed. Mice were briefly anesthetized under isofluroane, and all whiskers were trimmed on the object-presentation side of the face. Animals were returned to their home cage to fully recover from anesthesia. Then mice were placed into the behavioral apparatus and were exposed to the object discrimination task.

Chapter 4 : General Discussion

In this thesis, I have confirmed that barrel cortical activity is modulated by learning, with cells being recruited to represent a rewarded stimulus; meanwhile, repeated exposure to an unrewarded stimulus leads to habituation and reduces population of cells responding to the repeatedly presented stimulus. Furthermore, I have demonstrated that, over the course of learning (but not repeated exposure), barrel cortex gains the ability to encode the progression of time. Finally, I found that unexpected perturbations in the timecourse of trial events strongly drive activity in barrel cortex, potentially even more strongly than the tactile stimuli themselves.

4.1   Clocks and temporal coding schemes

As referenced in the Introduction (Chapter 1), a variety of brain regions have been hypothesized to constitute a "clock"—to track the passage of time or the duration of events. In addition, multiple coding schemes have been proposed as potential neural time-keeping mechanisms. One proposed method by which neurons may encode time is through individual cells' precise patterns of activity that fluctuate over the course of a trial or event. For instance, Bodner et al. (2005) identified putative "memory cells" in primate primary somatosensory cortex: these cells displayed activity patterns that spanned the duration of a delay in a working memory task (Bodner et al., 2005). While this delay-period activity does not necessarily explicitly track the duration of the delay or the time elapsed since

stimulus presentation, these types of temporal information might be contained in the neuronal activity patterns.

Another way in which cells might encode the time elapsed since an event is through a "ramping" neuronal responses, either as decay in stimulus-evoked responses—*i.e.*, "ramp-down"—or as a "ramp-up" to the time of a reward. In the "ramp-down" case, a long enough time constant of decay may make it possible to decode the amount of time after the presentation of a salient stimulus, based only on the instantaneous firing rate or response intensity of a cell or population of cells. Similarly, responses that "ramp-up" to the moment of a predictable event (Makino and Komiyama, 2015; Shuler and Bear, 2006) would similarly contain temporal information. However, one might expect the temporal resolution of this coding scheme to be rather low, as the slope of the ramps would have to be relatively shallow in order to encode information about longer time intervals. A steeper slope would allow for more accurate time encoding, but for shorter intervals.

Finally, a sequence-based coding scheme may be particularly effective (Zhou et al., 2020), whereby time-points are represented by distinct cells, which are active only at short intervals, and which respond at the same point of time in each trial. Hippocampal time cells operate in this way, with each cell active at a given moment in the trial. Similar sequences have been identified in striatum (Jin et al., 2009; Zhou et al., 2020) and prefrontal cortex (Jin et al., 2009). While time could not be decoded from individual cells in this coding scheme, the neuronal population could contain high-resolution temporal information.

It is possible that all of these mechanisms play a role in time-keeping across various brain regions, or even across different conditions within the same brain region. Specifically in my experiments, I have uncovered evidence for a sequence-based time-keeping approach in barrel cortex. Average responses of the cell population do appear to display sustained responses following a stimulus, which might initially suggest a ramp-down mechanism; yet, an examination of individual neurons reveals that cells respond in a very time-constrained manner, without any ramping. Instead, cells responses tile much of the trial duration: though many cells activate at particularly salient moments in the trial—*e.g.*, stimulus onset and offset—many cells exhibit time-constrained responses at moments throughout the trial, even when the stimulus is stationary, and indeed after it exits the whiskers' reach entirely.

4.2   Time coding and behavioral relevance

In our experiments, time was decodable only on rewarded, stimulus-present trials, not on stimulus-absent trials: if we trained the classifier on "hit" trials, we were able to decode time on both stimulus-present trial types ("hit" and "miss" trials). Time could not be extracted on "correct rejection" trials, regardless of the trial type used to train the decoder. One possible explanation is that, on stimulus-absent trials, somatosensory cues that denote the timecourse of the trial are not available to the barrel cortex. Indeed, while the sound of the motor that moves the wheel would likely provide the mouse with auditory information as to the progression of the trial, it is possible that cells of the barrel cortex do

not have access to this information. Perhaps one would be able to decode time on stimulus-absent trials from the auditory cortex, where the auditory timing cues are processed, and in higher-order brain regions where multi-sensory integration and explicit task representation is thought to take place. Nonetheless, in light of evidence that multi-sensory integration occurs even in primary sensory cortex, one might expect barrel cortex to have access to the auditory cues as well: specifically, there is evidence of auditory projections to the somatosensory cortex (Budinger et al., 2006) as well as auditory responses in somatosensory cortical cells (Maruyama and Komai, 2018; Zhou and Fuster, 2004). On the other hand, there is precedent for the observation that timing of stimuli of different modalities may not be encoded together in the same brain region (Bueti et al., 2008). Yet another possible explanation lies in the behavioral relevance of the trial types: unrewarded, stimulus-absent trials are perhaps uninteresting to the mouse, and as the mouse can simply ignore these trials, neural activity ceases to correlate with the passage of time.

Our behavioral paradigm is not well-suited to disambiguate these possibilities, nor can the experiments described above shed light on the relationship between time encoding and behavioral relevance. These questions may point to an interesting avenue for future research. Nevertheless, past work in non-sensory cortical regions and in subcortical structures provides a framework in which to investigate this issue. For instance, previous attempts to decode time in primate association cortical areas were only successful in instances where timing information was important for the task, and not when it was behaviorally irrelevant (Cueva et al., 2020). In dorsolateral striatum, on the other hand, Toso et al. (2021) were able to decode time even when the timing information was

uninformative to the animal (Toso et al., 2021b). Other studies agree that timing information need not be explicitly informative to allow for time coding, but do suggest that animals must be attentive to the task at hand (Jin et al., 2009). This latter finding supports my hypothesis that time cannot be decoded on unrewarded trials due to animals' inattentiveness on these trial types.

Our own behavioral paradigm can also be modified to determine whether behavioral relevance is important to time coding: for instance, rather than rewarding the object-present trials, we could deliver reward following the object-absent trials, making those trials more salient to the animal. If time coding emerges in this scenario, we would conclude events have to be behaviorally relevant for the progression of time to be encoded during those events; in addition, such a finding would suggest that time coding in a primary sensory region can depend on sensory information from a different modality.

4.3 Subjective time perception and its neural correlates

When does "subjective" time differ from veridical time? A classic example is the frequently-reported experience of "time dilation" during sudden, jarring, often life-threatening events. Such circumstances are difficult to experimentally induce and even harder to quantify—though some authors have tried (Stetson et al., 2007). However, less dramatic disturbances of time perception can be induced and measured in a controlled laboratory environment. From these experiments, a number of factors have been found to influence time perception: perceived stimulus duration varies with stimulus features such

as intensity, size, novelty, and numerosity (Toso et al., 2021a; Tse et al., 2004; Xuan et al., 2007). Those of us who have never been involved in life-threatening situations have almost certainly experienced temporal illusions and distortions in everyday life as well. Consider chronostasis, or the stopped-clock illusion, in which the first of a series of events or stimuli is perceived as lasting longer: during an initial glance at an analog clock, the second hand seems to stall for much longer than a second before it begins ticking at a normal pace. City-dwellers may experience a similar illusion while walking through busy streets: upon reaching a crosswalk, the pedestrian sees the blinking "don't walk" traffic-control beacon that indicates an imminent "don't walk" signal. Yet due to chronostasis, the blinking beacon appears constantly illuminated rather than blinking, to the point where the pedestrian may stop at the crosswalk rather than hurrying across[1].

While many examples of disturbed time perception exist, both in human and animal models, few studies have directly investigated the neural correlates of subjective time perception. But several have begun to uncover certain clues: using the aforementioned psychophysics evidence that stimulus intensity modulates perceived time duration, Toso et al. (2021b) found that striatal cells encode stimulus duration identically for high-intensity and low-intensity stimuli, which should have been perceived as relatively longer and shorter, respectively. In cortex, on the other hand, we might expect subjective time to be encoded, in light of a recent study that found that manipulations of cortical activity

---

[1] I am aware that the flashing phase of the signal is meant to convey to pedestrians to *finish* crossing but not to begin walking if they haven't already ventured into the crosswalk. Yet even the most risk-averse person cannot argue that the majority of a city population interprets the "flashing hand" as anything other than instruction to rush across the street.

affected both judgements of stimulus intensity *and* duration (Reinartz et al., 2021).

Meanwhile, in the cutaneous rabbit illusion (Geldard and Sherrick, 1972), which likely depends on temporal integration and postdiction, a signature of the illusory perception has been found in human S1 (Blankenburg et al., 2006).

There are several competing theories for the mechanism of the disturbances of time perception. From their experiments, Stetson et al., (2007) concluded that instances of "time dilation" do not involve any alternations in temporal resolution; rather, they suggest that the extra attentional allocation to frightening scenarios leads to richer memory of those events, in turn causing those events to be judged as lasting longer (Stetson et al., 2007). If this hypothesis is true, then disturbances of time perception are retrospective, and real-time perception of time progression is unchanged. This theory would explain why striatal cells were found to encode objective time, even if animals "report" that they perceived certain stimuli to last longer (Toso et al., 2021b). In addition, the study from Reinartz et al. (2021) could be interpreted in this light: stimulation of cortex did in fact alter perception of stimulus intensity; yet the corresponding changes in duration perception were simply retrospective judgements that came about as a result of what was effectively a manipulation of stimulus intensity. On the other hand, Salvioni et al. (2013) found that TMS interference with V1 activity modulated interval duration judgements, but did *not* affect the perception of nontemporal stimulus attributes.

In future experiments, computational methods of the kind I used to decode the progression of time from neural activity could test this hypothesis directly. If we find that objective time is encoded across the brain, then the retrospective, memory-based theory

likely explains time perception and its disturbances. However, if we find certain brain regions where cells encode subjective time—that is, if time dilation or compression is accompanied by a corresponding change in the slope of the diagonal line in a decoder confusion matrix (see Figure 2.11a)—that would suggest that the real-time experience of time progression, not just the memory of the event, is altered.

4.4   Conclusion: a distributed network for time-keeping

A significant novel finding of this thesis is that primary somatosensory cortex encodes time. Similar experiments in other modalities will likely yield similar findings in other primary sensory regions. Earlier in this thesis, I described temporal processing in a variety of brain regions—some of which have been posited to house the brain's "central clock" (see Introduction chapter). While temporal representations in various brain regions could be inherited from a "central clock" elsewhere, time and temporal expectation are increasingly thought to be encoded by a distributed network of brain areas, rather than localized in any given region (Bueti et al., 2010; Karmarkar and Buonomano, 2007). One piece of evidence for distributed time-keeping is that timing information about different modalities is represented in distinct regions (Bueti et al., 2008). Timing information may still have arrived at these distinct regions from elsewhere (*e.g.*, through feedback inputs from an upstream area), but the fact that the temporal representations differ between regions implies that they each may have a unique role in temporal processing. Future experiments to distinguish the representation of subjective versus objective time in

different brain areas may further test this idea. In the meantime, whether time-keeping arises due to intrinsic network properties in primary sensory cortex (Karmarkar and Buonomano, 2007), or whether timing information is carried to there from upstream regions and then molded to serve a behaviorally-relevant purpose, remains an open question.

Chapter 5 : Epilogue—or, thoughts on scientific subjectivity, controversy, and progress

To what degree is it acceptable for intuition to guide scientific exploration? It is likely unavoidable that we base our hypotheses at least partly on intuition—yet intuition is frequently wrong. On the other hand, I would argue that *introspection*, unlike intuition, should be embraced as a tool rather than frowned upon as a handicap. While intuition is an initial feeling, devoid of reasoning, introspection is contemplative and critical. In fact, introspection and contemplation may provide a way to overcome the shortcomings of our flawed intuition. And is not introspection a form of self-experimentation (if only in the form of a thought experiment)? Self-experimentation, though now frowned upon, has a long history in scientific exploration, and has led to important discoveries (though also occasionally ended with tragic consequences). With introspection, we can consider and mentally manipulate our own experiences without having to sacrifice ourselves for science.

After all, do we, as scientists, not marvel at all the idiosyncrasies of our strange existence? My interest in neuroscience was certainly built on a foundation of curiosity and confusion about the nature of my own experience. As many children do, I asked my parents questions about the self (*"why am I me and not someone else?"*), time and the continuity of experience (*"what changes as time passes? Am I still the same person I was a second ago? What if I sit really still and don't move at all, then will I stay the same?"*), and consciousness (*"do you see colors the same way I see colors? What is it like for my granddad to be*

colorblind?"). Years (decades?) later, I am completing my PhD in neuroscience, and I am still pondering these questions. I may have concluded (against all intuition) that the self is an illusion and that no distinction differentiates the brain and the mind. Yet the questions persist: *how does the brain generate a continuous subjective experience? How do we perceive the passage of time, and why is our perception of it so weird—with some moments dragging on interminably, while hours at a party (especially if the cocktails are plentiful) rush by?* These questions have propelled my PhD research about time coding, and drive my continuing interest in the neural correlates behind distortions of time perception.

*   *   *

In the last decade, I have learned, contemplated, often forgotten, and occasionally relearned a wide range of concepts, theories, and (maybe) facts. I have also accepted that science is at its core a human endeavor, inextricably tied to subjective intuition, and colored by emotion. Whatever objective reality exists must be inevitably interpreted through a framework of our own construction.

As such, it is natural that scientists begin to feel passionately about their own research and theories; when our ideas are unavoidably challenged, the ensuing defensive reaction is normal. But I have heard of instances where researchers tried to thwart new studies and prevent new findings from being published in an attempt to preserve the authority of their own ideas. Indeed, it is not uncommon among scientists to try to steer

clear of controversy when writing papers rather than to highlight discrepancies between studies, so as to not offend the authors cited. But controversy should be exciting!

If our goal is progress, we must be wary not to let expertise translate into dogma, nor to allow disagreement to slide toward animosity. We must be as open-minded as we are critical.

References

Ahmed, M.S., Priestley, J.B., Castro, A., Stefanini, F., Solis Canales, A.S., Balough, E.M.,
Lavoie, E., Mazzucato, L., Fusi, S., Losonczy, A., 2020. Hippocampal Network
Reorganization Underlies the Formation of a Temporal Association Memory. Neuron
107, 283-291.e6. https://doi.org/10.1016/j.neuron.2020.04.013

Alink, A., Blank, H., 2021. Can expectation suppression be explained by reduced attention
to predictable stimuli? NeuroImage 231, 117824.
https://doi.org/10.1016/j.neuroimage.2021.117824

Alink, A., Schwiedrzik, C.M., Kohler, A., Singer, W., Muckli, L., 2010. Stimulus
Predictability Reduces Responses in Primary Visual Cortex. J. Neurosci. 30, 2960–
2966. https://doi.org/10.1523/JNEUROSCI.3730-10.2010

Anderson, B., Peissig, J.J., Singer, J., Sheinberg, D.L., 2006. XOR style tasks for testing
visual object processing in monkeys. Vision Research 46, 1804–1815.
https://doi.org/10.1016/j.visres.2005.11.023

Arabzadeh, E., Panzeri, S., Diamond, M.E., 2006. Deciphering the Spike Train of a
Sensory Neuron: Counts and Temporal Patterns in the Rat Whisker Pathway. J.
Neurosci. 26, 9216–9226. https://doi.org/10.1523/JNEUROSCI.1491-06.2006

Ashe, J., Bushara, K., 2014. The olivo-cerebellar system as a neural clock. Adv Exp Med
Biol 829, 155–165. https://doi.org/10.1007/978-1-4939-1782-2_9

Avidan, G., Harel, M., Hendler, T., Ben-Bashat, D., Zohary, E., Malach, R., 2002. Contrast Sensitivity in Human Visual Areas and Its Relationship to Object Recognition. Journal of Neurophysiology 87, 3102–3116. https://doi.org/10.1152/jn.2002.87.6.3102

Ayaz, A., Stäuble, A., Hamada, M., Wulf, M.-A., Saleem, A.B., Helmchen, F., 2019. Layer-specific integration of locomotion and sensory information in mouse barrel cortex. Nat Commun 10, 2585. https://doi.org/10.1038/s41467-019-10564-8

Bakhurin, K.I., Goudar, V., Shobe, J.L., Claar, L.D., Buonomano, D.V., Masmanidis, S.C., 2017. Differential Encoding of Time by Prefrontal and Striatal Network Dynamics. J Neurosci 37, 854–870. https://doi.org/10.1523/JNEUROSCI.1789-16.2016

Bakin, J.S., Weinberger, N.M., 1990. Classical conditioning induces CS-specific receptive field plasticity in the auditory cortex of the guinea pig. Brain Res 536, 271–286. https://doi.org/10.1016/0006-8993(90)90035-a

Bale, M.R., Bitzidou, M., Giusto, E., Kinghorn, P., Maravall, M., 2020. Sequence Learning Induces Selectivity to Multiple Task Parameters in Mouse Somatosensory Cortex. Current Biology S096098222031602X. https://doi.org/10.1016/j.cub.2020.10.059

Banerjee, A., Parente, G., Teutsch, J., Lewis, C., Voigt, F.F., Helmchen, F., 2020. Value-guided remapping of sensory cortex by lateral orbitofrontal cortex. Nature 585, 245–250. https://doi.org/10.1038/s41586-020-2704-z

Barth, A.L., Poulet, J.F.A., 2012. Experimental evidence for sparse firing in the neocortex. Trends Neurosci 35, 345–355. https://doi.org/10.1016/j.tins.2012.03.008

Blankenburg, F., Ruff, C.C., Deichmann, R., Rees, G., Driver, J., 2006. The Cutaneous

　　Rabbit Illusion Affects Human Primary Sensory Cortex Somatotopically. PLOS

　　Biology 4, e69. https://doi.org/10.1371/journal.pbio.0040069

Bodner, M., Shafi, M., Zhou, Y.-D., Fuster, J.M., 2005. Patterned firing of parietal cells in a

　　haptic working memory task. European Journal of Neuroscience 21, 2538–2546.

　　https://doi.org/10.1111/j.1460-9568.2005.04085.x

Boynton, G.M., Demb, J.B., Glover, G.H., Heeger, D.J., 1999. Neuronal basis of contrast

　　discrimination. Vision Research 39, 257–269. https://doi.org/10.1016/S0042-

　　6989(98)00113-8

Brosch, M., Selezneva, E., Scheich, H., 2011. Representation of Reward Feedback in

　　Primate Auditory Cortex. Front. Syst. Neurosci. 5.

　　https://doi.org/10.3389/fnsys.2011.00005

Budinger, E., Heil, P., Hess, A., Scheich, H., 2006. Multisensory processing via early

　　cortical stages: Connections of the primary auditory cortical field with other sensory

　　systems. Neuroscience 143, 1065–1083.

　　https://doi.org/10.1016/j.neuroscience.2006.08.035

Bueti, D., Bahrami, B., Walsh, V., 2008. Sensory and Association Cortex in Time

　　Perception. Journal of Cognitive Neuroscience 20, 1054–1062.

　　https://doi.org/10.1162/jocn.2008.20060

Bueti, D., Bahrami, B., Walsh, V., Rees, G., 2010. Encoding of Temporal Probabilities in

　　the Human Brain. Journal of Neuroscience 30, 4343–4352.

　　https://doi.org/10.1523/JNEUROSCI.2254-09.2010

Chau, L.S., Akhtar, O., Mohan, V., Kondilis, A., Galvez, R., 2014. Rapid adult experience-dependent anatomical plasticity in layer IV of primary somatosensory cortex. Brain Res 1543, 93–100. https://doi.org/10.1016/j.brainres.2013.10.043

Chen, J.L., Margolis, D.J., Stankov, A., Sumanovski, L.T., Schneider, B.L., Helmchen, F., 2015. Pathway-specific reorganization of projection neurons in somatosensory cortex during learning. Nat Neurosci 18, 1101–1108. https://doi.org/10.1038/nn.4046

Chéreau, R., Bawa, T., Fodoulian, L., Carleton, A., Pagès, S., Holtmaat, A., 2020. Dynamic perceptual feature selectivity in primary somatosensory cortex upon reversal learning. Nat Commun 11, 3245. https://doi.org/10.1038/s41467-020-17005-x

Constantinople, C.M., Bruno, R.M., 2013. Deep Cortical Layers are Activated Directly by Thalamus. Science 340, 1591–1594. https://doi.org/10.1126/science.1236425

Cueva, C.J., Saez, A., Marcos, E., Genovesio, A., Jazayeri, M., Romo, R., Salzman, C.D., Shadlen, M.N., Fusi, S., 2020. Low-dimensional dynamics for working memory and time encoding. PNAS 117, 23021–23032. https://doi.org/10.1073/pnas.1915984117

David, S.V., Fritz, J.B., Shamma, S.A., 2012. Task reward structure shapes rapid receptive field plasticity in auditory cortex. Proc Natl Acad Sci U S A 109, 2144–2149. https://doi.org/10.1073/pnas.1117717109

den Ouden, H.E.M., Friston, K.J., Daw, N.D., McIntosh, A.R., Stephan, K.E., 2009. A Dual Role for Prediction Error in Associative Learning. Cerebral Cortex 19, 1175–1185. https://doi.org/10.1093/cercor/bhn161

Doherty, J.R., Rao, A., Mesulam, M.M., Nobre, A.C., 2005. Synergistic Effect of Combined Temporal and Spatial Expectations on Visual Attention. J. Neurosci. 25, 8259–8266. https://doi.org/10.1523/JNEUROSCI.1821-05.2005

Egger, R., Narayanan, R.T., Guest, J.M., Bast, A., Udvary, D., Messore, L.F., Das, S., de Kock, C.P.J., Oberlaender, M., 2020. Cortical Output Is Gated by Horizontally Projecting Neurons in the Deep Layers. Neuron 105, 122-137.e8. https://doi.org/10.1016/j.neuron.2019.10.011

Estebanez, L., Boustani, S.E., Destexhe, A., Shulz, D.E., 2012. Correlated input reveals coexisting coding schemes in a sensory cortex. Nat Neurosci 15, 1691–1699. https://doi.org/10.1038/nn.3258

Fiser, A., Mahringer, D., Oyibo, H.K., Petersen, A.V., Leinweber, M., Keller, G.B., 2016. Experience-dependent spatial expectations in mouse visual cortex. Nat Neurosci 19, 1658–1664. https://doi.org/10.1038/nn.4385

Fletcher, P.C., Frith, C.D., 2009. Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. Nat Rev Neurosci 10, 48–58. https://doi.org/10.1038/nrn2536

Fritz, J.B., David, S.V., Radtke-Schuller, S., Yin, P., Shamma, S.A., 2010. Adaptive, behaviorally-gated, persistent encoding of task-relevant auditory information in ferret frontal cortex. Nat Neurosci 13, 1011–1019. https://doi.org/10.1038/nn.2598

Gavornik, J.P., Bear, M.F., 2014. Learned spatiotemporal sequence recognition and prediction in primary visual cortex. Nat Neurosci 17, 732–737. https://doi.org/10.1038/nn.3683

Geldard, F.A., Sherrick, C.E., 1972. The Cutaneous "Rabbit": A Perceptual Illusion. Science. https://doi.org/10.1126/science.178.4057.178

Ghazanfar, A.A., Schroeder, C.E., 2006. Is neocortex essentially multisensory? Trends in Cognitive Sciences 10, 278–285. https://doi.org/10.1016/j.tics.2006.04.008

Gilbert, C.D., Sigman, M., 2007. Brain States: Top-Down Influences in Sensory Processing. Neuron 54, 677–696. https://doi.org/10.1016/j.neuron.2007.05.019

Goard, M., Dan, Y., 2009. Basal forebrain activation enhances cortical coding of natural scenes. Nat Neurosci 12, 1444–1449. https://doi.org/10.1038/nn.2402

Grill-Spector, K., Henson, R., Martin, A., 2006. Repetition and the brain: neural models of stimulus-specific effects. Trends in Cognitive Sciences 10, 14–23. https://doi.org/10.1016/j.tics.2005.11.006

Grill-Spector, K., Malach, R., 2004. The Human Visual Cortex. Annual Review of Neuroscience 27, 649–677. https://doi.org/10.1146/annurev.neuro.27.070203.144220

Hamm, J.P., Shymkiv, Y., Han, S., Yang, W., Yuste, R., 2021. Cortical ensembles selective for context. Proc Natl Acad Sci USA 118, e2026179118. https://doi.org/10.1073/pnas.2026179118

Henschke, J.U., Dylda, E., Katsanevaki, D., Dupuy, N., Currie, S.P., Amvrosiadis, T., Pakan, J.M.P., Rochefort, N.L., 2020. Reward Association Enhances Stimulus-Specific Representations in Primary Visual Cortex. Current Biology 30, 1866-1880.e5. https://doi.org/10.1016/j.cub.2020.03.018

Henson, R.N.A., 2003. Neuroimaging studies of priming. Progress in Neurobiology 70, 53–81. https://doi.org/10.1016/S0301-0082(03)00086-8

Herbst, S.K., Javadi, A.H., Meer, E. van der, Busch, N.A., 2013. How Long Depends on

    How Fast—Perceived Flicker Dilates Subjective Duration. PLOS ONE 8, e76074.

    https://doi.org/10.1371/journal.pone.0076074

Hong, Y.K., Lacefield, C.O., Rodgers, C.C., Bruno, R.M., 2018. Sensation, movement and

    learning in the absence of barrel cortex. Nature 561, 542–546.

    https://doi.org/10.1038/s41586-018-0527-y

Hutson, K.A., Masterton, R.B., 1986. The sensory contribution of a single vibrissa's

    cortical barrel. J Neurophysiol 56, 1196–1223. https://doi.org/10.1152/jn.1986.56.4.1196

Insanally, M.N., Carcea, I., Field, R.E., Rodgers, C.C., DePasquale, B., Rajan, K., DeWeese,

    M.R., Albanna, B.F., Froemke, R.C., 2019. Spike-timing-dependent ensemble

    encoding by non-classically responsive cortical neurons. eLife 8, e42409.

    https://doi.org/10.7554/eLife.42409

Jaramillo, S., Zador, A.M., 2011. The auditory cortex mediates the perceptual effects of

    acoustic temporal expectation. Nat Neurosci 14, 246–251.

    https://doi.org/10.1038/nn.2688

Jin, D.Z., Fujii, N., Graybiel, A.M., 2009. Neural representation of time in cortico-basal

    ganglia circuits. Proc Natl Acad Sci U S A 106, 19156–19161.

    https://doi.org/10.1073/pnas.0909881106

Karmarkar, U.R., Buonomano, D.V., 2007. Telling time in the absence of clocks. Neuron

    53, 427–438. https://doi.org/10.1016/j.neuron.2007.01.006

Kato, H.K., Gillet, S.N., Isaacson, J.S., 2015. Flexible Sensory Representations in Auditory

    Cortex Driven by Behavioral Relevance. Neuron 88, 1027–1039.

    https://doi.org/10.1016/j.neuron.2015.10.024

Keller, A.J., Houlton, R., Kampa, B.M., Lesica, N.A., Mrsic-Flogel, T.D., Keller, G.B.,

    Helmchen, F., 2017. Stimulus relevance modulates contrast adaptation in visual

    cortex. eLife 6, e21589. https://doi.org/10.7554/eLife.21589

Keller, G.B., Bonhoeffer, T., Hübener, M., 2012. Sensorimotor Mismatch Signals in

    Primary Visual Cortex of the Behaving Mouse. Neuron 74, 809–815.

    https://doi.org/10.1016/j.neuron.2012.03.040

Keller, G.B., Mrsic-Flogel, T.D., 2018. Predictive Processing: A Canonical Cortical

    Computation. Neuron 100, 424–435. https://doi.org/10.1016/j.neuron.2018.10.003

Kim, J., Erskine, A., Cheung, J.A., Hires, S.A., 2020. Behavioral and Neural Bases of Tactile

    Shape Discrimination Learning in Head-Fixed Mice. Neuron 108, 953-967.e8.

    https://doi.org/10.1016/j.neuron.2020.09.012

Kuhlman, S.J., O'Connor, D.H., Fox, K., Svoboda, K., 2014. Structural Plasticity within the

    Barrel Cortex during Initial Phases of Whisker-Dependent Learning. J. Neurosci. 34,

    6078–6083. https://doi.org/10.1523/JNEUROSCI.4919-12.2014

Labarrera, C., Deitcher, Y., Dudai, A., Weiner, B., Amichai, A.K., Zylbermann, N., London,

    M., 2018. Adrenergic Modulation Regulates the Dendritic Excitability of Layer 5

    Pyramidal Neurons In Vivo. Cell Reports 23, 1034–1044.

    https://doi.org/10.1016/j.celrep.2018.03.103

Lacefield, C.O., Pnevmatikakis, E.A., Paninski, L., Bruno, R.M., 2019. Reinforcement
Learning Recruits Somata and Apical Dendrites across Layers of Primary Sensory
Cortex. Cell Reports 26, 2000-2008.e2. https://doi.org/10.1016/j.celrep.2019.01.093

Lyamzin, D.R., Aoki, R., Abdolrahmani, M., Benucci, A., 2021. Probabilistic discrimination
of relative stimulus features in mice. PNAS 118.
https://doi.org/10.1073/pnas.2103952118

MacDonald, C.J., Lepage, K.Q., Eden, U.T., Eichenbaum, H., 2011. Hippocampal "time
cells" bridge the gap in memory for discontiguous events. Neuron 71, 737–749.
https://doi.org/10.1016/j.neuron.2011.07.012

Makino, H., Komiyama, T., 2015. Learning enhances the relative impact of top-down
processing in the visual cortex. Nat Neurosci 18, 1116–1122.
https://doi.org/10.1038/nn.4061

Margolis, D.J., Lütcke, H., Schulz, K., Haiss, F., Weber, B., Kügler, S., Hasan, M.T.,
Helmchen, F., 2012. Reorganization of cortical population activity imaged
throughout long-term sensory deprivation. Nat Neurosci 15, 1539–1546.
https://doi.org/10.1038/nn.3240

Martin, B., Wittmann, M., Franck, N., Cermolacce, M., Berna, F., Giersch, A., 2014.
Temporal structure of consciousness and minimal self in schizophrenia. Frontiers in
Psychology 5, 1175. https://doi.org/10.3389/fpsyg.2014.01175

Martins, A.R.O., Froemke, R.C., 2015. Coordinated forms of noradrenergic plasticity in
the locus coeruleus and primary auditory cortex. Nat Neurosci 18, 1483–1492.
https://doi.org/10.1038/nn.4090

Maruyama, A.T., Komai, S., 2018. Auditory-induced response in the primary sensory cortex of rodents. PLOS ONE 13, e0209266. https://doi.org/10.1371/journal.pone.0209266

Mather, J.G., Baker, R.R., 1981. Magnetic sense of direction in woodmice for route-based navigation. Nature 291, 152–155. https://doi.org/10.1038/291152a0

Mello, G.B.M., Soares, S., Paton, J.J., 2015. A Scalable Population Code for Time in the Striatum. Current Biology 25, 1113–1122. https://doi.org/10.1016/j.cub.2015.02.036

Mima, T., Nagamine, T., Nakamura, K., Shibasaki, H., 1998. Attention Modulates Both Primary and Second Somatosensory Cortical Activities in Humans: A Magnetoencephalographic Study. Journal of Neurophysiology 80, 2215–2221. https://doi.org/10.1152/jn.1998.80.4.2215

Morita, T., Kang, H., Wolfe, J., Jadhav, S.P., Feldman, D.E., 2011. Psychometric Curve and Behavioral Strategies for Whisker-Based Texture Discrimination in Rats. PLOS ONE 6, e20437. https://doi.org/10.1371/journal.pone.0020437

Nemec, P., Altmann, J., Marhold, S., Burda, H., Oelschlager, H.H., 2001. Neuroanatomy of magnetoreception: the superior colliculus involved in magnetic orientation in a mammal. Science 294, 366–368. https://doi.org/10.1126/science.1063351

Pachitariu, M., Stringer, C., Schröder, S., Dipoppa, M., Rossi, L.F., Carandini, M., Harris, K., 2016. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. https://doi.org/10.1101/061507

Pantoja, J., Ribeiro, S., Wiest, M., Soares, E., Gervasoni, D., Lemos, N.A.M., Nicolelis, M.A.L., 2007. Neuronal Activity in the Primary Somatosensory Thalamocortical Loop

Is Modulated by Reward Contingency during Tactile Discrimination. Journal of Neuroscience 27, 10608–10620. https://doi.org/10.1523/JNEUROSCI.5279-06.2007

Park, J.M., Hong, Y.K., Rodgers, C.C., Dahan, J.B., Harano, N., Schmidt, E.R., Bruno, R.M., 2022. Deep and superficial layers of the primary somatosensory cortex are critical for whisker-based texture discrimination in mice. https://doi.org/10.1101/2020.08.12.245381

Pastalkova, E., Itskov, V., Amarasingham, A., Buzsáki, G., 2008. Internally generated cell assembly sequences in the rat hippocampus. Science 321, 1322–1327. https://doi.org/10.1126/science.1159775

Peron, S.P., Freeman, J., Iyer, V., Guo, C., Svoboda, K., 2015. A Cellular Resolution Map of Barrel Cortex Activity during Tactile Behavior. Neuron 86, 783–799. https://doi.org/10.1016/j.neuron.2015.03.027

Pinto, L., Goard, M.J., Estandian, D., Xu, M., Kwan, A.C., Lee, S.-H., Harrison, T.C., Feng, G., Dan, Y., 2013. Fast modulation of visual perception by basal forebrain cholinergic neurons. Nat Neurosci 16, 1857–1863. https://doi.org/10.1038/nn.3552

Pitas, A., Albarracín, A.L., Molano-Mazón, M., Maravall, M., 2017. Variable Temporal Integration of Stimulus Patterns in the Mouse Barrel Cortex. Cerebral Cortex 27, 1758–1764. https://doi.org/10.1093/cercor/bhw006

Pleger, B., Blankenburg, F., Ruff, C.C., Driver, J., Dolan, R.J., 2008. Reward Facilitates Tactile Judgments and Modulates Hemodynamic Responses in Human Primary Somatosensory Cortex. J Neurosci 28, 8161–8168. https://doi.org/10.1523/JNEUROSCI.1093-08.2008

Pleger, B., Ruff, C.C., Blankenburg, F., Klöppel, S., Driver, J., Dolan, R.J., 2009. Influence of Dopaminergically Mediated Reward on Somatosensory Decision-Making. PLoS Biol 7, e1000164. https://doi.org/10.1371/journal.pbio.1000164

Pluta, S., Naka, A., Veit, J., Telian, G., Yao, L., Hakim, R., Taylor, D., Adesnik, H., 2015. A direct translaminar inhibitory circuit tunes cortical output. Nat Neurosci 18, 1631–1640. https://doi.org/10.1038/nn.4123

Pluta, S.R., Telian, G.I., Naka, A., Adesnik, H., 2019. Superficial layers suppress the deep layers to fine tune cortical coding. J. Neurosci. https://doi.org/10.1523/JNEUROSCI.1459-18.2018

Poort, J., Khan, A.G., Pachitariu, M., Nemri, A., Orsolic, I., Krupic, J., Bauza, M., Sahani, M., Keller, G.B., Mrsic-Flogel, T.D., Hofer, S.B., 2015. Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. Neuron 86, 1478–1490. https://doi.org/10.1016/j.neuron.2015.05.037

Poort, J., Wilmes, K.A., Blot, A., Chadwick, A., Sahani, M., Clopath, C., Mrsic-Flogel, T.D., Hofer, S.B., Khan, A.G., 2022. Learning and attention increase visual response selectivity through distinct mechanisms. Neuron 110, 686-697.e6. https://doi.org/10.1016/j.neuron.2021.11.016

Ramirez, A., Pnevmatikakis, E.A., Merel, J., Paninski, L., Miller, K.D., Bruno, R.M., 2014. Spatiotemporal receptive fields of barrel cortex revealed by reverse correlation of synaptic input. Nat Neurosci 17, 866–875. https://doi.org/10.1038/nn.3720

Ramirez, J.J., Colwill, R.M., 2012. Pavlovian Biconditional Discrimination Learning in the C57BL/6J Mouse. Behav Processes 90, 278–286. https://doi.org/10.1016/j.beproc.2012.02.017

Reinartz, S., Fassihi, A., Paz, L., Pulecchi, F., Gigante, M., Diamond, M.E., 2021. The sensory code within sense of time. https://doi.org/10.1101/2021.07.29.454157

Richter, D., Ekman, M., de Lange, F.P., 2018. Suppressed Sensory Response to Predictable Object Stimuli throughout the Ventral Visual Stream. J Neurosci 38, 7452–7461. https://doi.org/10.1523/JNEUROSCI.3421-17.2018

Rodgers, C.C., DeWeese, M.R., 2014. Neural correlates of task switching in prefrontal cortex and primary auditory cortex in a novel stimulus selection task for rodents. Neuron 82, 1157–1170. https://doi.org/10.1016/j.neuron.2014.04.031

Rodgers, C.C., Nogueira, R., Pil, B.C., Greeman, E.A., Park, J.M., Hong, Y.K., Fusi, S., Bruno, R.M., 2021. Sensorimotor strategies and neuronal representations for shape discrimination. Neuron 109, 2308-2325.e10. https://doi.org/10.1016/j.neuron.2021.05.019

Salvioni et al. - 2013 - How the Visual Brain Encodes and Keeps Track of Ti.pdf, n.d.

Salvioni, P., Murray, M.M., Kalmbach, L., Bueti, D., 2013. How the Visual Brain Encodes and Keeps Track of Time. Journal of Neuroscience 33, 12423–12429. https://doi.org/10.1523/JNEUROSCI.5146-12.2013

Schwark, H.D., Malpeli, J.G., Weyand, T.G., Lee, C., 1986. Cat area 17. II. Response properties of infragranular layer neurons in the absence of supragranular layer

activity. Journal of Neurophysiology 56, 1074–1087.

https://doi.org/10.1152/jn.1986.56.4.1074

Shimbo, A., Izawa, E.-I., Fujisawa, S., 2021. Scalable representation of time in the

hippocampus. Science Advances 7, eabd7013. https://doi.org/10.1126/sciadv.abd7013

Shuler, M.G., Bear, M.F., 2006. Reward timing in the primary visual cortex. Science 311,

1606–1609. https://doi.org/10.1126/science.1123513

Skaggs, W., McNaughton, B., Gothard, K., 1993. An Information-Theoretic Approach to

Deciphering the Hippocampal Code, in: Advances in Neural Information Processing

Systems. Morgan-Kaufmann.

Stetson, C., Fiesta, M.P., Eagleman, D.M., 2007. Does Time Really Slow Down during a

Frightening Event? PLOS ONE 2, e1295.

https://doi.org/10.1371/journal.pone.0001295

Summerfield, C., Trittschuh, E.H., Monti, J.M., Mesulam, M.-M., Egner, T., 2008. Neural

repetition suppression reflects fulfilled perceptual expectations. Nat Neurosci 11,

1004–1006. https://doi.org/10.1038/nn.2163

Teufel, C., Fletcher, P.C., 2020. Forms of prediction in the nervous system. Nat Rev

Neurosci 21, 231–242. https://doi.org/10.1038/s41583-020-0275-5

Toso, A., Fassihi, A., Paz, L., Pulecchi, F., Diamond, M.E., 2021a. A sensory integration

account for time perception. PLoS Comput Biol 17, e1008668.

https://doi.org/10.1371/journal.pcbi.1008668

Toso, A., Reinartz, S., Pulecchi, F., Diamond, M.E., 2021b. Time coding in rat dorsolateral

striatum. Neuron. https://doi.org/10.1016/j.neuron.2021.08.020

Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J.J., Moser, M.-B., Moser, E.I., 2018. Integrating time from experience in the lateral entorhinal cortex. Nature 561, 57–62. https://doi.org/10.1038/s41586-018-0459-6

Tse, P.U., Intriligator, J., Rivest, J., Cavanagh, P., 2004. Attention and the subjective expansion of time. Perception & Psychophysics 66, 1171–1189. https://doi.org/10.3758/BF03196844

Umbach, G., Kantak, P., Jacobs, J., Kahana, M., Pfeiffer, B.E., Sperling, M., Lega, B., 2020. Time cells in the human hippocampus and entorhinal cortex support episodic memory. PNAS 117, 28463–28474. https://doi.org/10.1073/pnas.2013250117

Weis, T., Brechmann, A., Puschmann, S., Thiel, C.M., 2013. Feedback that confirms reward expectation triggers auditory cortex activity. Journal of Neurophysiology 110, 1860–1868. https://doi.org/10.1152/jn.00128.2013

Woods, A.J., Hamilton, R.H., Kranjec, A., Minhaus, P., Bikson, M., Yu, J., Chatterjee, A., 2014. Space, Time, and Causality in the Human Brain. Neuroimage 92, 285–297. https://doi.org/10.1016/j.neuroimage.2014.02.015

Xin, Y., Zhong, L., Zhang, Y., Zhou, T., Pan, J., Xu, N., 2019. Sensory-to-Category Transformation via Dynamic Reorganization of Ensemble Structures in Mouse Auditory Cortex. Neuron 103, 909-921.e6. https://doi.org/10.1016/j.neuron.2019.06.004

Xuan, B., Zhang, D., He, S., Chen, X., 2007. Larger stimuli are judged to last longer. Journal of Vision 7, 2. https://doi.org/10.1167/7.10.2

Zhang, M., Kwon, S.E., Ben-Johny, M., O'Connor, D.H., Issa, J.B., 2020. Spectral hallmark of auditory-tactile interactions in the mouse somatosensory cortex. Communications Biology 3, 1–17. https://doi.org/10.1038/s42003-020-0788-5

Zhou, S., Masmanidis, S.C., Buonomano, D.V., 2020. Neural Sequences as an Optimal Dynamical Regime for the Readout of Time. Neuron 108, 651-658.e5. https://doi.org/10.1016/j.neuron.2020.08.020

Zhou, Y.-D., Fuster, J.M., 2004. Somatosensory cell response to an auditory cue in a haptic memory task. Behav Brain Res 153, 573–578. https://doi.org/10.1016/j.bbr.2003.12.024

Zmarz, P., Keller, G.B., 2016. Mismatch Receptive Fields in Mouse Visual Cortex. Neuron 92, 766–772. https://doi.org/10.1016/j.neuron.2016.09.057

Zuo, Y., Perkon, I., Diamond, M.E., 2011. Whisking and whisker kinematics during a texture classification task. Philos Trans R Soc Lond B Biol Sci 366, 3058–3069. https://doi.org/10.1098/rstb.2011.0161