A FUNCTIONAL CORTICAL NETWORK FOR SENSORIMOTOR SEQUENCE GENERATION

by
Duo Xu

A dissertation submitted to Johns Hopkins University in conformity with
the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland
May 2020

Abstract

The brain generates complex sequences of movements that can be flexibly reconfigured in real-time based on sensory feedback, but how this occurs is not fully understood. We developed a novel 'sequence licking' task in which mice directed their tongue to a target that moved through a series of locations. Mice could rapidly reconfigure the sequence online based on tactile feedback. Closed-loop optogenetics and electrophysiology revealed that tongue/jaw regions of somatosensory (S1TJ) and motor (M1TJ) cortex encoded and controlled tongue kinematics at the level of individual licks. Tongue premotor (anterolateral motor, ALM) cortex encoded intended tongue angle in a smooth manner that spanned individual licks and even whole sequences, and progress toward the reward that marked successful sequence execution. ALM activity regulated sequence initiation, but multiple cortical areas collectively controlled termination of licking. Our results define a functional cortical network for hierarchical control of sensory- and reward-guided orofacial sequence generation.

Advisor and Primary Reader

    Daniel H. O'Connor, Ph.D.

Secondary Reader

    Kristina J. Nielsen, Ph.D.

Thesis Committee

    Paul F. Worley, Ph.D.

    Kristina J. Nielsen, Ph.D.

    Solange P. Brown, Ph.D.

    Shreesh P. Mysore, Ph.D.

Acknowledgements

unusual mixed into his signature smile, and once done rushed back home to take his sick kid to hospital. The harmonious and supportive lab environment comes from both the friendly and intelligent members and Dan's artful leadership. I started to appreciate that more after mentoring students on my own. I am grateful to Yuxi Chen, Angel Delgado, Natasha Hughes, Ali Aly and others who trusted my good intentions and lived with my strong (frequently biased) opinions. They showed amazing professionalism, diligence and energy bandwidth already at their young age. It is both exciting and fearful to imagine one day, I would have to compete with people like them.

If our lab is like a family, the Hopkins neuroscience community would be my extended one. The intellectual support and challenges from people around me have deeply shaped the way I think. More importantly, people made me feel at home here. There was not much of a cultural shock. I was not worried too much about my broken English especially when I first came to the U.S.. Silly enough that I once wanted to experience "racism", only because I was so protected in the Hopkins' shield and became disconnected with the weight of that word. Illusion or not, I never thought Baltimore is a dangerous city and I still do not.

More than what can be expressed in words, I want to thank my previous mentor Dr. Muming Poo at the Institute of Neuroscience in Shanghai. It was not a total coincidence that the thesis is about sequences. His inspiration on neural sequences seeded my neuroscience worldview, although it was not until much later did I fully appreciate its significance. Muming is one of his kind, making amazing contributions from one field to another. He said "it is just one of many different styles of research", but it is a style that strongly resonates with me, though having the ability to do that may be another thing that "can only happen but not be asked for."

During the six years in the U.S., I only went back to China once. I was absent from my dear grandma's funeral, after my mom's serious injury, and during my dad's difficult time. Being almost thirty years old, I still acted like a teenager who habitually ignored parents and tended to do the opposite of what they would wish for. Science is my haven, but I sometimes

use it for escape from responsibilities. And the less my parents asked, the more I wanted to escape. Their never changing love to me and the lack of complaint only made it worse. My mom always liked the motto of my primary school: "to be a person everyone likes, a person everyone needs". However, I have had many excuses not to be a person everyone likes. I had believed that everyone has opinions about each other, and saying them out indiscriminately is a way to be self-consistent and honest. I also had aversion against certain types of conversation, which distanced myself from friendly people. Now, I felt maybe I had turned everything upside down. Perhaps the *goal* after all is not to hold some moral rules (as arbitrary and self-prescribed as they can be), but simply "to be a person everyone likes." Life is neither meaningful or meaningless within an individual because one cannot define what meaning even means; meaning has always been defined by *others*. In the end, I want to thank everyone that has made me a better self, and I sincerely apologize to those I have hurt and disappointed.

Contents

List of Figures

Chapter 1

Introduction

A baby is born. She is crying, kicking, knowing nothing about this strange world. But the world does not appear in total randomness (Simoncelli and Olshausen, 2001). A falling toy will not suddenly freeze in the air, float back up, or liquify; where she moves her hand toward at this moment must be where she moves her hand away from next; a meaningful phrase she hears is not a string of arbitrary sounds, and her first utterance of "ma-ma" is rewarded by warm laughters of parents. Grasping patterns like these is a natural instinct since we are born (Saffran and Kirkham, 2018) and is the foundation to develop advanced intelligence throughout our life, from exploring new environments to mastering motor skills and perhaps to building an internal model of the world which we can imagine within. Channeling the physical world, our senses transmit continuous streams of signals that are related in space and time. What often concurs with these sensory streams is a sequence of motor signals. The motor sequence changes the sensory streams and the sensory updates in turn inform future motor actions, forming an intricate sensorimotor loop (Hayhoe, 2017; Gallivan et al., 2018). Relationships embedded in these sensorimotor sequences hold the secrets of how the world and us together evolve.

Studies of sensorimotor sequences span different functional and anatomical domains of the nervous system. On one end, specialized circuits, such as those found in the spinal cord

and brainstem responsible for stereotypic walking (Kiehn, 2016), breathing and chewing (Moore et al., 2014), govern detailed sequencing of muscle activations with specific integration of sensory or homeostatic feedback. Due to genetically determined wiring and limited plasticity, these circuits trade functional flexibility with reliability, accuracy and readiness for use. On the other end, highly plastic circuits in cerebrum and cerebellum have greater potential to adapt for changing demands guided by sensory experiences and objectives (Doya, 2000; Hikosaka et al., 2002). Most skillful behaviors involve supervised and memorized sensorimotor sequences (Desrochers et al., 2016). These sequences often exhibit a hierarchical organization where elementary patterns of sensory input and motor kinematics form short sequences that are structured into longer and more complex compositions (Rosenbaum, 2009; Zhang and Ghazanfar, 2020). Though memorized, these sequences must be flexibly rearranged or branched upon arrival of relevant sensory input (Gallivan et al., 2018). A fascinating and unresolved question, therefore, is how the brain organizes such hierarchical representation and flexible control.

This thesis aims (1) to establish a novel behavioral paradigm in mice which marries skillful, hierarchical and flexible sequence behaviors with a wide selection of experimental tools to measure and manipulate the nervous system, (2) to explore the level of control and flexibility that mice can achieve with coordinated tongue-jaw movements, (3) to define the cortical regions that are important for this sequence performance, and (4) to characterize and understand the neural representations across a functional network that underlies the behavioral sequence generation. It also describes technical innovations arising from these investigations that have broader applications.

1.1   Behavioral Sequencing

To put the features and utilities of the novel task into context, we first need to ask what is unique about sequence behaviors, and why are they worth studying? We will then review

classic behavioral models in human and non-human species that have provided windows to look into the neural underpinning of the sequence behaviors, consider each of their strengths and limitations, and in the end point out the gaps filled and the opportunities offered by the present task. Afterall, a model that does not address meaningful questions would be useless no matter how amenable it is to genetic manipulations or how many advanced tools have been developed for it.

1.1.1   Characteristics of Behavioral Sequences

One may be tempted to reduce a sequence of behaviors to a chain of simple stimulus-response associations, where the sensory state caused by a previous response serves as a stimulus to trigger the next response. Despite meaningful contributions (Clower, 1998), this kind of response chaining theories have been largely rejected by both theoretical considerations and experimental evidence (Rosenbaum, 2009). The main point of this section therefore is not to further criticize these theories, but to highlight the key ingredients sequence behaviors entail beyond isolated sensorimotor transformation.

Sequences with diverging states cannot be achieved via chaining a series of fixed sensorimotor transformations. When we write, type or pronounce the two words, "sequence" and "sequencing", the initial part "sequen" are the same, which produces the same sensory states, whereas the ending has to differ, which is not possible given any fixed stimulus-response pair. Common slips of the tongue, such as a spoonerism of saying "kitchen" as "chicken", suggest that a later sensorimotor component is already available early on. Context-dependent sequence generation is not unique to speech, and has been widely studied in the planning and production of saccade sequences (Hayhoe and Ballard, 2005) and object manipulation (Cohen and Rosenbaum, 2004). In general, informative sensory cues are not always present when a motor program needs to be selected or unfold differently. The brain therefore must keep a memory of past states, guide ongoing movements, and plan for future actions by inference (Niv, 2019). Despite the computational distinction between movements that can

be causally chained and those requiring inference, the nervous system would likely solve the infinite number of real life tasks by the same sets of circuits and via universal mechanisms. We will expand on that later when reviewing surprising neural representations found even in simple sequences.

Another problem of response chaining theories is their inability to explain the generalization from one behavioral pattern to another. Every language learner would have the experience that saying a new phrase can be hard but substituting different words becomes much easier once the syntax becomes familiar. Motor sequences learned with one hand can facilitate, or transfer to, the performance using the other hand (Shea et al., 2011). Both children and macaques are able to learn a multi-touch sequence and generalize to perform its reversed version (Jiang et al., 2018). Generalization can be accounted for by a hierarchical organization of behavioral elements as it offers control at different levels of abstraction (Rosenbaum, 2009). For instance, phonemes are organized in words, which are in turn organized in phrases, sentences, etc., which can be flexibly reorganized to form novel expressions according to grammatical rules. A motor sequence can be decomposed into subsequences, or chunks, which are separated by extended time gaps and/or increased error rate from one chunk to another (Rosenbaum et al., 1986; Jin and Costa, 2015). Sequence chunking is not only a result of experimental quantification but also an intrinsic property of the sensorimotor system necessary for proficient performance (Sakai et al., 2003). Learning of a new sequence involves formation of chunks and gradual shortening of gaps within and between chunks (Jin and Costa, 2015).

A single response to a cue has a natural end point - when the response is completed, it is over. However, for motor sequences, especially those with periodic movements such as walking, a stopping mechanism would be necessary to terminate the ongoing movements based on internal state or external cue (Fujii and Graybiel, 2003). In a study about motor output buffering, typists in the middle of a typing sequence were able to stop promptly in response to an auditory cue, emitting only one extra key press on average (Logan, 1982). Pa-

4

tients with Parkinson's and Huntington's disease, however, exhibited deficits in terminating a movement sequence (Benecke et al., 1987; Agostino et al., 1992).

In sum, a behavioral sequence is composed of a series of elementary movements, usually organized in a hierarchy; planning and memory are often necessary to determine the course of a sequence in addition to sensory cues; a sequence can be executed and terminated automatically and independent from sensory feedback, or flexibly based on specific sensory feedback and internal states. In addition, timing each movement is also an important aspect of sequence behaviors.

1.1.2 Behavioral Models

Human sequence behaviors can be dazzlingly complex, hard to control experimentally, and not feasible for other animals to learn. When developing a behavioral task, it is an art to distill the essential elements of sequencing while keeping the form simple. Tanji and colleagues trained macaques to hold a handle and wait for a start signal to initiate three movements of pushing, pulling and turning in a specific sequence in order to obtain a reward at the end (Tanji and Shima, 1994). We will later refer this as the Push-Pull-Turn task. There are two nice features about this design. First, since there is no external cue (in one version of this task) about what movement should be executed next, monkeys have to keep the sequence in mind before and during the execution. This memory-guided motor sequencing differs from other tasks where each movement is guided by a distinct sensory cue. Second, consecutive movements are separated by a short interval when the hand is passively returned to a neutral position. This sets the initial sensory state of the motor plant identical before monkeys execute the next movement. As a result, movement sequencing cannot be achieved by simply learning a chain of mapping from the last sensory state to the next movement. Instead, a memory about sequence context (e.g. serial order) must be maintained. A major limitation of this task is the difficulty to study hierarchical sequences due to the limited number of possible movement combinations. Hikosaka and colleagues developed a task which

requires monkeys or humans to perform a two-level hierarchy of motor sequences (Hikosaka et al., 1995; Sakai et al., 2003). There is a 4-by-4 array of buttons for a monkey to press. A lower level sequence, or a subsequence, consists of pressing two simultaneously lit buttons in the correct order. A series of five randomly determined subsequences comprise a higher level sequence, thus the name "2 × 5 task" (or "2 × 10 task" for humans). This task is able to produce a huge number of arbitrary sequences. Although subjects were not asked to remember the order of subsequences in a sequence, they tend to spontaneously predict the upcoming subsequences with anticipatory movements before the light comes on. In another word, this task, once well-learned, is naturally performed in a memory-guided fashion. One may notice that tasks like these tend to ignore, or intentionally control for, the influence of the sensory inputs during sequence performance. This helps to distinguish a sensorimotor association process from a memory-guided one. However, they do not offer an opportunity to study the integration of sensory input with memory and how inference can be made.

Saccades inherently incorporate sensory and motor components and usually occur in sequences (Hayhoe and Ballard, 2005). An average human induces 100-thousand saccades each day, making it perhaps our most frequently performed sensorimotor sequence behavior (unless someone is really talkative). In most saccade sequence tasks, subjects learned to make directional saccades sequentially to a set of targets. For example, the Pull-Push-Turn and "2 × 5" tasks both have variations that replace hand movements with directional saccades (Isoda and Tanji, 2002; Lu et al., 2002; Berdyyeva and Olson, 2010). Other tasks use more natural behaviors such as sequences of free-viewing scans for a hidden target (Desrochers et al., 2015a).

Human speech, in my opinion, is the pinnacle of sensorimotor sequence control. In non-human species, vocalization can also be a sophisticated sensorimotor sequence behavior. When songbirds sing, the vocal organs need to be orchestrated with rich dynamics and millisecond precision. Songs are acquired by juvenile birds imitating the calls of adult tutor birds with a growing vocabulary and an evolving repertoire (Okubo et al., 2015). There

is a clear hierarchical organization of motor elements in zebra finch songs. A typical bout of song is composed of one or more motifs, each of which consists of a rapid sequence of 3-8 syllables, and the syllables may be further divided into notes (Brown, 2017). As a beautiful and intricate behavior, birdsong is worth to be understood in its own right. It also provides unique opportunities to study how sequences are shaped at developmental timescale (Brainard and Doupe, 2013). From a primate-centric point of view, however, it is not clear to what extent we can find generalizable mechanisms from such a highly specialized system. Perhaps, a fundamental difference between birdsongs and the learned sequence behaviors in primates is the way behavioral goals (or loss functions) are defined, which may lead to very different neural architectures.

Like humans and non-human primates, rodents exhibit a myriad of goal-oriented sequence behaviors. Navigating through space is inherently a sensorimotor sequence, which often involves a synthesis of memory and sensory cues to guide the decision at each turn and corner. Studies of spatial navigation in rodents (Foster, 2017) and relational representation in general (Wilson et al., 2014; Miller et al., 2017; Aronov et al., 2017; Liu et al., 2019) have led to a tremendous breakthrough in our understanding of how the brain organizes related states of the world. However, the brain has multiple processing and learning systems. Mechanisms found in the hippocampal-prefrontal system may not explain procedural behaviors which rely more on the sensorimotor system (McDonald and Hong, 2013), though the boundary can be blurry.

Rodents do not have a long history of use in studying refined sensorimotor control. Many of the new studies began by adapting primate behavioral tasks. For example, mice have been recently trained to move a cursor to target via a joystick (Panigrahi et al., 2015), to adapt to force fields when operating a manipulandum (Mathis et al., 2017), or to reach for a food pellet (Guo et al., 2015) or water droplet (Galiñanes et al., 2018). Although more sophisticated than most rodent tasks, these behaviors do not involve sensorimotor sequences. Costa and colleagues developed a motor sequence task where mice learned to

press a lever repetitively for a fixed number of times before getting a reward (Jin and Costa, 2010). There is only one fixed sequence to be learned which is composed of a single type of movement (i.e. lever pressing) with unconstrained time intervals between movements. Due to its minimalist design, this task has been limited to mainly study sequence initiation and termination. However, even with relatively crude behaviors and limited applications, mouse models allow one to attribute neural representations to defined circuit components more easily and to apply better controlled neural perturbations, thanks to an ever growing arsenal of genetic tools, the amenity to new electrophysiological and imaging techniques, and a fast experimental turnover for testing hypotheses. But can it be better?

We can only study the neural basis of a brain function if it is present in a task performed by a given species. A useful rodent task for sensorimotor sequence research should therefore exhibit features unique to sequence behaviors, such as hierarchy, memory-guidance, and flexible control. However, it was not known whether or not rodents are capable of behaviors with such complexity. The task developed in this thesis demonstrates that mice can indeed master memory-guided motor sequences using their tongues, rapidly reconfigure the sequence online based on tactile feedback, and string sequences into a higher level composition. It fills the gap and provides new opportunities to inform generalizable mechanisms for sensorimotor sequence control.

1.2   Neural Basis of Behavioral Sequences

A behavior can be explained at different levels - psychological, representational, architectural (circuit), cellular and molecular. If neuroscience follows the path of other natural sciences, our understanding deepens when we are able to reduce an upper level explanation to the next. In the past three decades, cognitive and systems neuroscientists have made significant progress to bridge psychological observations of sequence behavior to neural representations and, to certain extent, circuit architectures (Georgopoulos, 2000; Desrochers et al., 2016;

Mackevicius and Fee, 2018). This section reviews how different brain regions and neurons encode and control motor sequence execution, how expected and unexpected sensory sequences are processed, and theoretical models for sensorimotor sequence processing.

1.2.1 Coding and Control of Motor Sequences

The first links between sequence behavior and its neural substrates came from clinical studies. Patients with lesions in supplementary motor area (SMA) and certain surrounding regions showed deficits when performing alternating movements with two hands (Laplane et al., 1977), sequential limb movements (Dick et al., 1986), speech (Goldberg, 1985), and temporally controlled finger movements (Halsband et al., 1993). Damage to the supplementary eye field (SEF) caused poor performance in sequential saccades, especially when the patients had to rely on memory in the absence of visual guidance (Pierrot-Deseilligny et al., 1995). The prefrontal cortex (PFC) is also involved in motor sequencing (Kolb and Milner, 1981; Kimura, 1982). Patients with Parkinson's (Harrington and Haaland, 1991; Agostino et al., 1992; Martin et al., 1994) and Huntington's (Thompson et al., 1988) diseases showed impairments in motor sequence tasks, suggesting a role for basal ganglia in sequence control.

In search of brain activities specific to sequence execution or learning, early human functional imaging studies usually took advantage of the unique features of complex motor sequence mentioned before, and contrasted that with simple movements. In these studies (Roland et al., 1980; Shibasaki et al., 1993; Jenkins et al., 1994; Kami et al., 1995; Catalan et al., 1998; Grosbras et al., 2001), many brain regions were found to be differentially activated in sequence performance or learning, which include not only SMA, anterior cingulate cortex (ACC), lateral premotor cortex (LPM), lateral PFC, the parietal cortex, and parts of striatum, but also primary motor (M1) and primary sensory (S1) cortex. However, it was not entirely clear to what extent the increased activity in each area was due merely to a greater demand for motor drive or to signals for sequencing (Tanji, 2001).

Over the years, several brain regions involved in sensorimotor sequence behaviors have
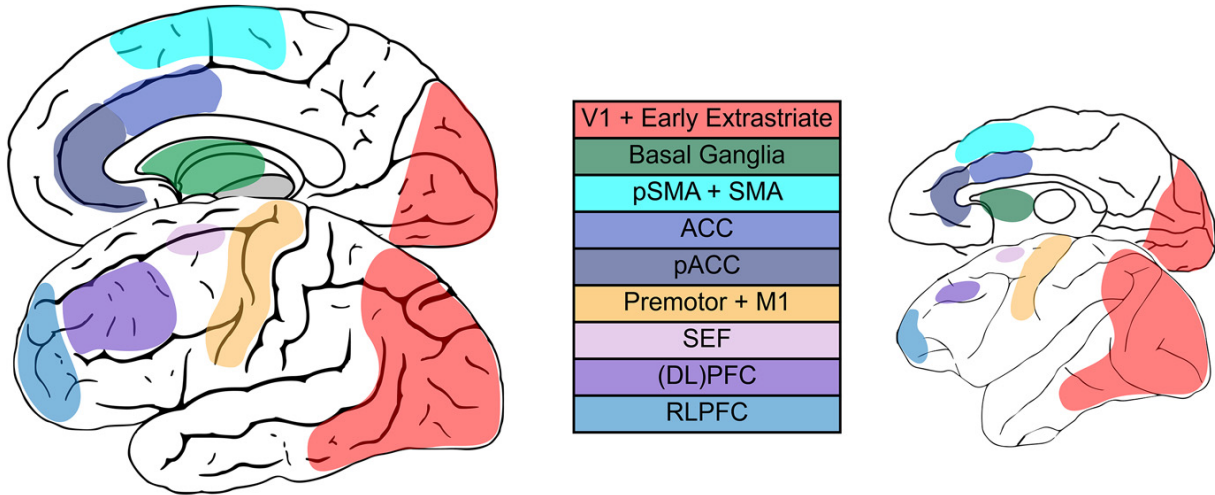
Figure 1.1: Cortical regions related to motor sequence control

From Desrochers et al. (2016). "Representation of main areas involved in sequential control in the human (left) and monkey (right) brain."

been characterized using single-unit recording and pharmacological inactivation in non-human primates (Figure 1.1). Tanji and colleagues compared neuronal coding from M1 and SMA in the Push-Pull-Turn task (Tanji and Shima (1994); Figure 1.2). They found that M1 neurons only responded for a specific movement regardless of which sequence it was in. In SMA, however, 26% of neurons fired in anticipation of a specific sequence, and another 36% of neurons that fired during the sequence responded to a specific movement at a certain rank order (e.g. a Pull as the second movement). M1 neurons also did not distinguish sensory- versus memory-guided movements whereas SMA neurons were more involved in memory-guided sequences and PM neurons more in sensory-guided ones (Mushiake et al., 1991). Traditionally defined SMA was divided into a rostral pre-SMA and SMA proper. Unlike SMA, 64% of neurons in pre-SMA became active when monkeys were learning a new sequence (Shima et al., 1996; Clower and Alexander, 1998). Consistent with these electrophysiological results, pharmacological inactivation of SMA and pre-SMA did not impair simple arm reaching movements but caused inability to perform memory-guided, but not visually-guided, reaching sequences (Shima and Tanji, 1998). Inactivations of pre-SMA, but
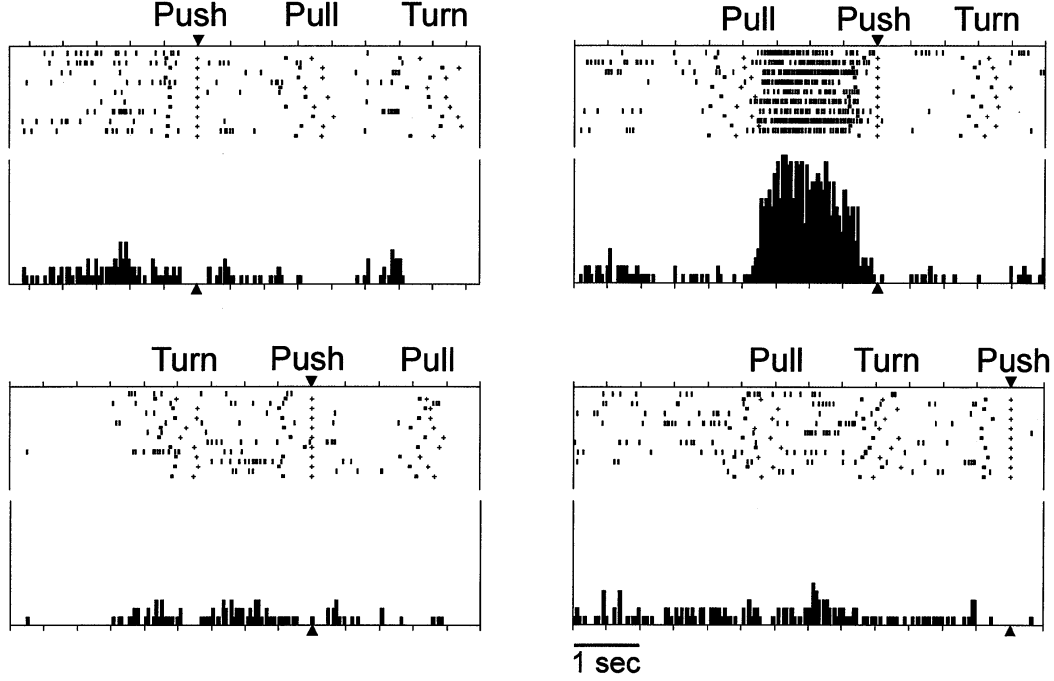
Figure 1.2: Sequence coding in supplementary motor area

From Tanji (2001) "Discharges of a supplementary motor area neuron whose activity increased selectively during the interval after completion of a particular movement, PULL, and before the initiation of another particular movement, PUSH. In raster displays, each row represents a trial, and dots represent individual discharges of this neuron. Small crosses denote the time of occurrence of the movement onset. In histograms, discharges over 12 trials are summated. Bin width for the display purpose is 40 ms. Triangles at the top of each raster indicate the start of the first movement."

not SMA, have been shown to cause deficits in learning new sequences in the "2 × 5" task (Hikosaka et al., 1995). In addition to phasic responses to individual movements, SMA neurons showed slow ramping activities throughout the course of a movement sequence when monkeys performed a hand cycling task (Russo et al., 2019). It will be demonstrated in this thesis that neurons in the anterolateral motor (ALM) cortex of mice also exhibit a dominant "ramping" activity, or more precisely, a smoothly changing activity that tracks the progression of a sequence.

Supplementary eye field (SEF) is an area close to SMA in the rostrolateral direction. In a saccadic version of the "2 × 5" (Lu et al., 2002) and Pull-Push-Turn (Isoda and Tanji, 2002; Berdyyeva and Olson, 2010) task, SEF neurons were found to encode movements in a

context-dependent manner similar to SMA neurons for arm and hand sequences. Intracortical microstimulation (ICMS) interfered with the sequencing of saccades and caused deficits in saccade kinematics (Berdyyeva and Olson, 2014). Despite the name, SEF may not be the only region for sequencing eye movements, as ICMS in SMA was also able to evoke saccades (Mushiake et al., 1991) and activities related to sequential saccades have also been found in SMA (Clower and Alexander, 1998; Berdyyeva and Olson, 2010).

Further rostromedial from pre-SMA is ACC, where most neurons showed selectivity to rank order and were activated more when monkeys were learning new sequence orders through trial-and-error (Procyk et al., 2000), reminiscent of pre-SMA coding but seeming to have more emphasis on rank order. Although the majority of studies of ACC have been focused on value representation, model updating and outcome evaluation, there could be deeper connections to sequence behaviors in terms of behavioral flexibility and persistence (Kolling et al., 2016).

Barone and Joseph were the first to observe single neurons coding for sequence structures in the brain (Barone and Joseph, 1989). They recorded neurons from dorsolateral prefrontal cortex (dlPFC) when monkeys learned to keep a sequence of three visual targets in working memory and then, in the absence of ordinal cues, reproduce that sequence by reaching to those targets in the correct order. About 12% of neurons maintained a working memory-like activity tuned to a specific target direction at a particular rank. During sequence execution, another 32% of neurons responded to movement toward a specific target direction, but only did so given certain subsequent movements but not others. Interestingly, neuronal coding across SMA, pre-SMA, SEF and dlPFC recorded in a same saccade task all showed rank selectivity without any qualitative difference (Berdyyeva and Olson, 2010).

Things we do in daily life are usually not isolated. Instead, some are grouped together to serve a small goal, and multiple small goals together achieve a bigger one (Lashley, 1951). This kind of cognitive hierarchy can be different from typical movement hierarchy because each component of the former is a goal-oriented sub-task. PFC hosts a cascade of executive

processes from PM to anterior PFC that controls behaviors at levels from sensory input to perceptual context to temporal episode (Koechlin et al., 2003). Desrochers et al. found that the BOLD response in rostrolateral PFC (rlPFC), but not pre-PMd, ramped as participants progressed through a sequence of tasks. Transcranial magnetic stimulation (TMS) to rlPFC, but not pre-PMd, resulted in compromised task performance (Desrochers et al., 2015b).

The cerebral cortex works together with basal ganglia and the cerebellum with topographically matched recurrent loops (Parent and Hazrati, 1995; Hunnicutt et al., 2016). It is not surprising to find neural activities in these subcortical structures related to sequence behaviors, and to see behavioral deficits when the activities are perturbed. In the "2 × 5" task, inactivation of middle and posterior putamen impaired execution of well-learned sequences, whereas inactivating anterior caudate and putamen only affected the learning of new sequences (Miyachi et al., 1997). Reducing cholinergic (Van Den Bercken and Cools, 1982) or dopaminergic (Matsumoto et al., 1999) transmission in striatum also caused deficits in sequence behaviors. After a movement sequence becomes habitual, striatal neurons were found to produce phasic response at the onset and offset of the sequence (Jin and Costa, 2010; Jin et al., 2014; Desrochers et al., 2015a).

The human and non-human primate models are in some sense double-edged swords. On one hand, we cannot fully understand human intelligence without studying them, since rodents simply lack the counterparts to certain specialized brain regions that are important for our unique cognitive functions (Preuss, 1995). On the other hand, with many of these specialized regions in primates showing similar (and sometimes ambiguous) activities and contributions, it could be useful to study qualitatively similar behaviors using the "simplified" rodent brains where the general circuit architecture is very relatable. Additionally, being able to reduce the representational level explanation to the circuit level will further help to elucidate the underlying mechanisms. Understanding rodent brain functions also has its own significance. For example, is there a network of brain regions that produces memory-guided yet flexible sensorimotor sequences? What could be the coding principles, circuit

13

architectures, and computational tradeoffs that allow smaller brains to pack these and other functions into its less parcelated and limited real estate?

1.2.2 Representations of Sensory Sequences

In healthy individuals, every motor output produces sensory feedback. A sequence of motor outputs naturally leads to a sequence of sensory inputs. Being able to recognize and utilize specific sensory patterns is key to behavioral adaptability and construction of internal models. For example, if the sensory sequence deviates from an expected one, detecting the "surprises" can inform the motor system to steer its course of action. In contrast to the classical sensory prediction error, sequence recognition and error detection may also depend on its sensory or motor contexts. The difference is analogous to how M1 and SMA neurons can both respond to a specific movement while the latter's activity also depends on abstract sequence structures (Tanji and Shima, 1994).

Even in the absence of motor output, the brain is able to passively pick up spatial and temporal sensory patterns. In anesthetized (or awake) rats, a light dot sweeping from one place to another in the visual field could elicit sequential firing of neurons in the primary visual cortex (V1). After repeated presentation of such sweeps, a single flash of light at the starting point recalled sequential firing in those neurons as if a sweeping dot were displayed (Xu et al., 2012). Repeated presentation of a grating sequence could lead to potentiation of V1 response in a sequence-specific manner where an altered element would induce a weaker response (Gavornik and Bear, 2014). Oftentimes, expected sensory sequences result from self-generated movements. As mice ran through a linear track with different visual patterns on the wall, a subset of V1 neurons were found to signal unexpected changes in the otherwise familiar visual flow (Keller et al., 2012; Fiser et al., 2016; Zmarz and Keller, 2016). A later study found that A24b (a subdivision of ACC) and the adjacent secondary motor cortex (M2) provided prediction signals to V1 (Leinweber et al., 2017). Similar effects of sensory cancellation by self-motions and prediction error signals were also observed in the auditory

14

cortex (Schneider et al., 2018).

The most ubiquitous feedback to motor output is tactile and proprioceptive. It is surprising to see how little we know about their sequence representation in the brain. Among the limited literature, an fMRI study asked subjects to discriminate if two tactile sequences are identical, and found activations in S1, SMA and dlPFC (Numminen et al., 2004). A recent study characterized neuronal tuning to sequential stimulation of adjacent whiskers in mouse primary somatosensory cortex (S1) (Laboy-Juárez et al., 2019), although it did not examine contextual level coding. Interestingly, when mice used whiskers to navigate through a tactile augmented reality, neurons in the whisker S1 did not tend to distinguish self-motion induced or decoupled sensory inputs (Sofroniew et al., 2015). The increasing evidence (Schneider, 2020) of sequence related signals in other modalities and the partially contrasting results from the whisker system expose a big gap in somatosensory research on sensory sequence representation. Although sequence level sensory prediction and prediction error is out of the scope of this thesis, the novel behavioral paradigm provides a great opportunity to further explore these topics.

1.2.3   Theories and Models

It has been argued that even if we knew "every action potential from every neuron"[1] and the entire connectome of the brain, we would still not be able to understand how the nervous system works. Understanding (including misunderstanding) does not come from data, but from models.

Models speak in different levels of abstraction. Here, we will discuss example models which reduce psychological/behavioral level explanation to representational level, and/or from representational to architectural/circuit level. Models also have different scopes. Some, for instance, may be concerned only with neural signal generation; others may cover the entire

---

[1] Erin McKiernan.   Apr 2, 2013.   President Obama's brain map project is hardly the next Human Genome.   *The Guardian.*   `https://www.theguardian.com/commentisfree/2013/apr/02/president-obama-brain-mapping-project-not-ideal`

sensorimotor loop plus learning mechanisms.

Recurrent neural networks (RNN) sit in the spotlight for sequence recognition and generation (LeCun et al., 2015). The simplest RNN consists of a group (or layer) of units that transforms activity from the input side to the output side, and the same output also loops back to the units themselves, usually with a time delay. This configuration allows the network to retain and accumulate information from past "selves" across time, and use this memory, in combination with new input, to generate the output. Via supervised, unsupervised or reinforcement learning, "synaptic" weights are adjusted and the network will produce desired sequences of activity based partly on the external (or "sensory") input and partly on the internal memory.

The behaviors of a RNN in the context of motor control can be understood from a dynamical systems point of view (Shenoy et al., 2013). Each instance of the RNN output describes a state living in the state-space of all possible states. A dynamical system specifies the rules of change where the next state is a function of the current state, external inputs, and noise (Strogatz, 2001). Once an appropriate initial state is set, a sequence of states will be iteratively rolled out along a trajectory. For memory-guided motor sequences, different trajectories must occupy different domains of the state-space, especially in the presence of noise, otherwise progression along one trajectory can be accidentally shunted to another, causing unwanted behaviors to occur. Differential preparatory activities and various the movement-rank selective responses (Barone and Joseph, 1989; Tanji and Shima, 1994; Clower and Alexander, 1998; Lu et al., 2002; Berdyyeva and Olson, 2010) can be useful to maintain unique evolution of population activity for different sequences. A recent work has shown that both artificial RNN and biological neurons in M1 have large separation, or little entanglement, among different state trajectories (Russo et al., 2018).

Although RNNs can carry out the core computation of the dynamics, having RNNs alone is not enough to account for real life behaviors. That would require more integrated and specialized networks which can handle complex sensorimotor transformation, representation

of internal states, and learning mechanisms (Richards et al., 2019). For example, Banino et al. connected two modules of LSTM (a variant of RNN), two modules of convolutional neural networks (CNN), several fully connected layers, reward signals, and etc. together to model how an individual (e.g. rat) learns to perform spatial navigation based on visual and vestibular feedback (Banino et al., 2018). They found that place and grid cell-like responses naturally emerged in specific modules of this network as a result of exploration. Another RNN based model deals specifically with hierarchical motor sequences (Yamashita and Tani, 2008). It has three layers of RNNs with fast to slower temporal dynamics. Proprioceptive and tactile inputs and motor outputs are all connected with the fastest RNN. After the model being trained on an hierarchical object manipulation task, units in different layers show hierarchical representations of the task structure. The faster layer signals detailed movements whereas the slower layer signals contexts of movement sequences. The activity in the slower layer also keeps state trajectories of different tasks separate despite the shared individual movements.

1.3   The Control of Tongue

Sensorimotor control of the tongue is involved in a wide range of both innate and learned behaviors, some essential for survival, such as communication and feeding, others mind-boggling, like knotting a cherry stem[2]. Mechanosensitive primary afferents transduct tactile information from the surface of the tongue to brainstem. From there, the main ascending pathway goes to the VPM thalamus and then to the somatosensory cortex. Sensorimotor transformation occurs across multiple sensory, motor, and cognitive areas (Crochet et al., 2019). To produce movements, motor commands are sent to the brainstem and then to the hypoglossal nucleus where motor neurons control the eight muscles in our tongue. This section will first introduce the tongue as a motor plant, then move on to the brainstem circuits

---

[2]Malia Wollan. June 23, 2017. How to Knot a Cherry Stem With Your Tongue. *The New York Times Magazine*, https://www.nytimes.com/2017/06/23/magazine/how-to-knot-a-cherry-stem-with-your-tongue.html

that generate innate patterns of tongue movement, and in the end review the neural basis of goal-oriented tasks where licking has been used as a major form of behavioral readout.

1.3.1 The Tongue

Mammalian tongues are classified into Types I and II (Doran, 1975). Type II tongues are heavily used for extra-oral exploration such as food gathering (Monotremata, Marsupilia, Pholidota) and can protrude more than 100% of its resting length. Type I tongues mainly function inside the mouth such as during eating, and can protrude less than 50% of their resting length. Both being Type I, the human and mouse tongues share more similarities than differences (Treuting et al., 2017). They have roughly the same shape, though the two halves are more obvious in mice. Both tongues are based near epiglottis and extend to the incisors. The dorsal surface is roughened with papilla and the ventral side is smooth (though still keratinized in mice). However, mice lack a lingual frenulum, the vertical fold of mucous membrane that connects the human tongue to the floor of the mouth. With less attachment between the anterior tongue to the mouth floor, mice are able to protrude the tongue a bit longer (proportionally) than humans can.

The tongue is controlled by four extrinsic (genioglossus, hyoglossus, styloglossus, and palatoglossus) muscles that go from bones into the tongue, and four intrinsic (superior longitudinal, inferior longitudinal, vertical, and transverse) muscles that live inside the tongue (McClung and Goldberg, 2000; Aoyagi et al., 2015). The incompressibility of liquid gives the tongue a fixed total volume. When the tongue is constricted circumferentially, it will extend longitudinally; when the volume of the left half shrinks, the right expands. Therefore, the intrinsic muscles can change the shape of the tongue by altering the distribution of stress from within. The extrinsic muscles can move the tongue forward or backward, and pull it to the left or right, in coordination with intrinsic muscles. Retrograde tracing showed that motor neurons in different subdivisions (though with overlap) of the hypoglossal nucleus innervate different tongue muscles (McClung and Goldberg, 1999). Although sixteen muscles

(eight on each side) are involved in controlling the tongue motion, it is not clear if the tongue has sixteen degrees of freedom since it depends on whether or not different muscles can be activated independently by upstream circuits.

1.3.2 Brainstem Circuits

The brainstem hosts a complex network of nuclei and circuits that funnel outputs into hypoglossal nuclei to drive tongue movements (Figure 1.3). There are three major themes in these control circuits - rhythm, reflex, and descending control.

Many behaviors are rhythmic in nature, such as breathing and chewing. In rodents, whisking and licking are typically also rhythmic. Although seemingly unrelated, these behaviors tend to oscillate at a similar (or harmonic) frequency in a phase-locked fashion (Kurnikova et al., 2017). Behind the scene, the preBotzinger complex is regarded as the master oscillator that shares a common oscillatory signal to many clients. In addition to initiating breathing, this complex resets the premotor oscillator for whisking (Moore et al., 2013; Deschênes et al., 2016) and is thought to serve a similar function in other rhythmic orofacial movements including licking and vocalization (Kleinfeld et al., 2015). However, the rhythm of tongue motions can shift from the rhythm of breathing (Welzl and Bureš, 1977) to

Figure 1.3: Diagram of brainstem circuits for tongue control

From McElvain et al. (2018). "'Feedforward' circuit diagram for muscles of the tongue. Arrow signifies the direction of signal flow; red signifies inhibitory and black signifies excitatory connections. Compiled data as described in the text. Many feedback connections and interconnections among premotor structures have been excluded for simplicity. Dashed lines are connections based on functional rather than anatomical data. Abbreviations: PrV (principal trigeminal nucleus); SpVO (spinal subnucleus oralis); SpVC (spinal subnucleus caudalis); Mes-V (mesencephalic sensory nucleus); VPMdm (ventral posterior medial nucleus, dorsomedial); VPMpc (ventral posterior medial nucleus, parvocellular division); S1 (primary sensory) cortex; ALM (anterior lateral motor) cortex; SNr (substantia nigra pars reticulata); CeA (central amygdala); Pontine/ cerebellum (circuit from pontine nuclei through cerebellar deep nuclei); lateral hypothal (lateral hypothalamus); PPN (pedunculopontine nucleus); Raphe (Raphe nuclei); PreBotC (preBotzinger respiratory complex) hIRt (hypoglossal intermediate reticular zone); IRt (intermediate reticular formation); Peri-V (peri-trigeminal area); PCRt (parvocellular reticular formation); dMRF (dorsal midbrain reticular formation); LPGi (lateral paragigantocellular reticular formation); and MdD (dorsal medullary reticular formation)."
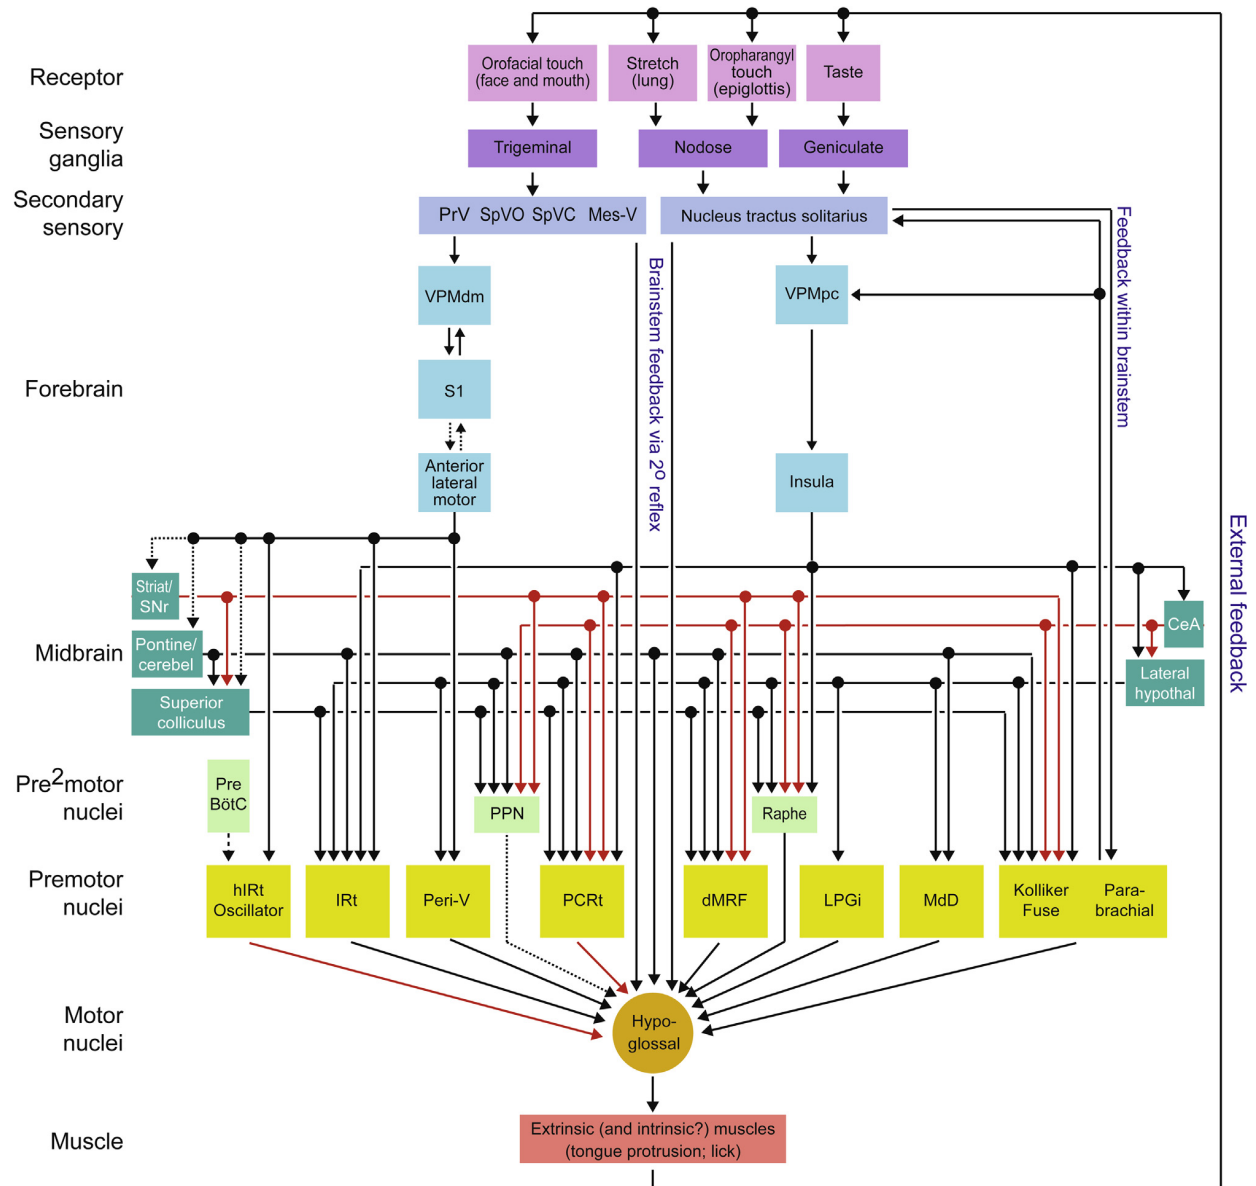
Figure 1.3: Diagram of brainstem circuits for tongue control

that of chewing (Travers et al., 2010). A hypoglossal subregion of intermediate reticular formation (hIRt) has been suggested as part of a separate licking oscillator (Travers et al., 1997; Stanek et al., 2014). So far, the circuit mechanisms for the muscular sequencing of rhythmic licking and the coordination with other orofacial structures are still largely unknown.

There is an abundance of reflexes that the tongue participates in (Miller, 2002) and many of the underlying circuits have been characterized (McElvain et al., 2018). For example, touching the tongue can orient it toward the stimulus (Bosma, 1973). Pulling the jaw open increases tongue protrusion (Miller, 2002). Infusing water into the mouth of anesthetized rats could induce reflexive swallowing (Grill and Norgren, 1978). As many learned sensorimotor sequences are composed of a hierarchy of simpler units of movement, it is important to know which require descending control, which can be mediated automatically via brainstem circuits, and under what circumstances the dependency can be changed.

Motor cortex is the most extensively studied region for direct descending control of tongue in primates (Sessle, 2011), rodents (Svoboda and Li, 2018) (Figure 1.4), and other species. Pyramidal tract neurons in the anterolateral motor area (ALM), a subregion of the secondary motor cortex (M2) in mice, send direct projections to the hIRt (Li et al., 2015). Intracortical microstimulation and optogenetic activation of ALM can induce rhythmic licking-like tongue and jaw movements (Komiyama et al., 2010; Li et al., 2015; Allen et al., 2017), whereas inhibition can change aspects of licking kinematics (Bollu et al., 2019). The ventrolateral striatum receives afferents from orofacial cortex (Hunnicutt et al., 2016) and the basal ganglia output neurons in substantia nigra pars reticulata project to various brainstem premotor regions (McElvain et al., 2018). Like in the cortex, stimulating these regions can evoke orofacial movements, though with abnormal forms (Delfs and Kelley, 1990; Inchul et al., 2005). The Crus I and II of cerebellum and the fastigial nucleus have been suggested to mediate tongue related sensorimotor processing (Bryant et al., 2010; Gaffield and Christie, 2017; Lu et al., 2013; Gao et al., 2018). Neurons in fastigial nuclei project directly to the hypoglossal nuclei (Stanek et al., 2014). Stimulation of different deep cerebellar nuclei was
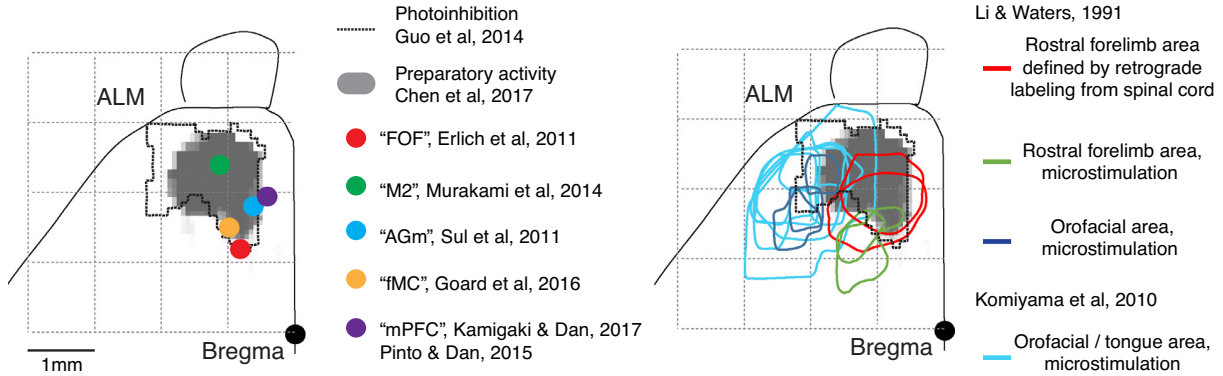
Figure 1.4: Functional mapping of the rodent motor cortex

Adapted from Svoboda and Li (2018). Region in the motor cortex with various recording locations (left) and microstimulation locations (right) superposed. The outline corresponds to the standard mouse brain from Allen Mouse Common Coordinate Framework. Note that studies on rats (FOF, M2) were scaled (scaling factor, 1.6, rat:mouse) to account for differences in brain size.

able to evoke different tongue movements (Bowman and Aldes, 1980).

1.3.3 Cortical Mechanisms in Goal-oriented Behaviors

Goal-directed behaviors differ from reflex-like behaviors in that they are executed voluntarily and are shaped by rewards and punishments. In operant tasks, subjects are trained to perform goal-directed behaviors by making appropriate motor responses to sensory cues. This process involves multiple subsystems in the central nervous system (Doya, 2000; Hikosaka et al., 2002), including various regions in the cerebral cortex that participate in the sensorimotor transformation.

The cortical mechanisms of sensory processing and motor control of the tongue were first studied in non-human primates. The tongue and jaw subregions in primate M1 and S1 are situated at the lateral ends of pre and postcentral sulcus, respectively. Monkeys were trained in a tongue protrusion task where they needed to use their tongue to push a button with a specific amount of force, and in a biting task where the biting force needed to be controlled (Murray and Sessle, 1992c). Many M1 and S1 neurons were modulated during one or both

tasks (Murray and Sessle, 1992b; Lin et al., 1994). In the tongue protrusion task, temporary bilateral cooling of M1 made monkeys unable to exert enough force for a sufficient period of time (Murray et al., 1991). Cooling of S1 increased the variability of tongue protrusion force both within and across trials (Lin et al., 1993). However, neither M1 or S1 cooling had an effect in the biting task, suggesting a difference in cortical involvement. In a modified task where monkeys performed tongue protrusions in three directions, 45% of the neurons in M1 tongue subregion showed directional selectivity in both preparatory and executive periods (Murray and Sessle, 1992a). Simultaneous recording in orofacial S1 and M1 showed strong directional coding in both regions but with different amplitudes and temporal profiles (Arce et al., 2013). It is, however, hard to investigate the circuit basis for the coordination of these two regions in monkeys.

More recently, a lot of progress has been made on the circuit mechanisms for sensorimotor transformation, motor planning, and working memory using goal-oriented licking tasks in mice (Svoboda and Li, 2018). In one such task, a vertical bar is presented to a mouse within its whisker field. The mouse needs to find, or "sample", the location of the bar by whisking against it and, after a delay without the bar, is required to make a left or right lick based on the previously sampled bar position (Guo et al., 2014). Mice not only need to acquire the correct sensorimotor mapping between different bar locations to lick directions, but also to hold the information in working memory between the sampling and response periods. A systematic screening via optogenetic inhibition across the dorsal cortex defined whisker S1 as a site for processing the tactile information during sampling, and ALM as a key region for holding a preparatory activity of upcoming lick direction during the delay period and later sending out the motor command. Detailed characterization of sensory and motor responses suggested a flow of sensorimotor transformation from whisker M2 to ALM, across ALM cortical layers (Chen et al., 2017), and across pyramidal tract neuron subtypes (Economo et al., 2018). ALM forms strong recurrent connections between the two hemispheres (Li et al., 2016), with motor thalamus (Guo et al., 2017), and indirectly with cerebellum (Gao et al.,

2018; Chabrol et al., 2019). This network is robust to perturbation in part of the circuit and may support an attractor dynamics for discrete actions (Inagaki et al., 2019). M1TJ neurons have also been found to encode lick direction in an operant task (Mayrhofer et al., 2019). Given similar neuronal tuning in M1TJ and ALM, further studies are needed to examine the coding and functional differences of these two regions, perhaps in more complicated tasks.

How similar mouse ALM and M1 are to the primate motor cortices (such as M1, PM, and SMA) is not clear. Although a one-to-one match does not seem possible, we do find similar logic of corticofugal projections, and signals for movements, motor planning, and possibly reward expectation (Chabrol et al., 2019). These perhaps partly reflect a related evolutionary path and partly reflect similar computational demands.

Although most of the tasks in mice discussed so far use tongue movements (typically licking) as a behavioral readout, the sensory input or feedback from the tongue is rarely studied. The tongue is able to receive somatosensation and taste at the same time. Mice have been trained to perform operant tasks with cues of different tastes, and various task-correlated neuronal activities in the gustatory cortex (GC) and orbitofrontal cortex (recognized as secondary GC) were reported (MacDonald et al., 2009; Yoshida and Katz, 2011; Fonseca et al., 2018). Tactile (and thermal) information from the tongue is represented in S1TJ (Clemens et al., 2018) although it is unclear if they can also be found in GC. During active sensing, the motor output is naturally linked with the sensory feedback (Graham et al., 2014; Severson et al., 2017). However, there has not been any study looking into how the S1TJ processes tactile feedback in a behavioral task.

1.4   Main Objectives

To sum up, generating complex and flexible sensorimotor sequences is at the core of advanced brain functions. Our understanding about the underlying neural mechanisms will continue to deepen as new studies bridge the gap between sequence behaviors and their circuit basis.

Different behavioral models in different species tend to balance a tradeoff among behavioral complexity, tractability and neural accessibility. With advances in technologies of recording and manipulating neural circuits, greatly improved tools for behavioral measurements, and the observations of non-primate species performing more complex tasks, a mouse model for sensorimotor sequence behavior has become imaginable. Like humans, mice exhibit rich orofacial behaviors. However, little is known about how sensorimotor sequences of the tongue can be generated in the brain.

In Chapter 2, we establish a novel 'sequence licking' task in which mice directed their tongue to a target that moved through a series of locations. We show that mice could rapidly reconfigure the sequence online based on tactile feedback.

In Chapter 3, via a systematic optoinhibition screening, we identify that ALM and M1TJ mediate sequence initiation and maintain the magnitude of licking kinematics during execution; S1TJ is important to guide the lick in the correct sequence; multiple anterior cortical regions are responsible for active sequence termination.

In Chapter 4, we first characterize the single-unit response properties and how they temporally tile the sequence performance. At population level, we reveal that S1TJ, M1TJ, and ALM neurons hierarchically encode behavioral variables with increasing levels of abstraction.

In Chapter 5, we zoom in on two types of high level coding by ALM neurons. First, the neuronal activities for sequence progress are modulated by reward signals. Second, ALM neurons maintain a working memory across trials which suggests a mechanism for composing a sequence of sequences.

In Chapter 6, we will summarize and expand on our findings, discuss limitations of the present study, and propose future directions.

1.5 General Methods

1.5.1 Mice

All procedures were in accordance with protocols approved by the Johns Hopkins University Animal Care and Use Committee. Mice were kept in a reverse light-dark cycle with each phase lasting 12 hours. Prior to surgery, mice were housed in groups of up to 5, but afterwards housed individually. Ten mice (8 male, 2 female) were obtained by crossing VGAT-IRES-Cre (Jackson Labs: 028862; B6J.129S6(FVB)-Slc32a1$^{\text{tm2(cre)Lowl}}$/MwarJ; Vong et al. (2011)) with Ai32 (Jackson Labs: 012569; B6;129S-Gt(ROSA)26Sor$^{\text{tm32(CAG-COP4*H134R/EYFP)Hze}}$/J) (Madisen et al., 2012) lines. Two (1 male, 1 female) were heterozygous VGAT-ChR2-EYFP (Jackson Labs: 014548; B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP)8Gfng/J) (Zhao et al., 2011) mice. Eight (6 male, 2 female) were wild-type mice, including seven C57BL/6J (Jackson Labs: 000664) mice and one wild-type littermate of the VGAT-ChR2-EYFP mice. Two were male TH-Cre 1 (Jackson Labs: 008601; B6.Cg-7630403G23Rik$^{\text{Tg(Th-cre)1Tmd}}$/J) (Savitt, 2005) mice. Two (1 male, 1 female) were Advillin-Cre (Jackson Labs: 032536; B6.129P2-Avil$^{\text{tm2(cre)Fawa}}$/J) (Zhou et al., 2010) mice. Mice ranged in age from ∼2-9 months at the start of training. A set of behavioural testing sessions typically lasted ∼1 month.

1.5.2 Surgery

Prior to behavioural testing, mice underwent the implantation of a metal headpost. For surgical procedures, mice were anesthetized with isoflurane (1-2%) and kept on a heating blanket (Harvard Apparatus). Lidocaine was used as a local analgesic and injected under the scalp at the start of surgery. Ketoprofen was injected intraperitoneally to reduce inflammation. All skin and periosteum above the dorsal surface of the skull was removed. The temporal muscle was detached from the lateral edges of the skull on either side and the bone ridge at the temporal-parietal junction was thinned using a dental drill to create a wider accessible region. Metabond (C & B Metabond) was used to cover the entirety of the skull

surface in a thin layer, seal the skin at the edges, and cement the headpost onto the skull over the lambda suture.

To make the skull transparent, a thin layer of cyanoacrylate adhesive was then dropped over the entirety of the Metabond-coated skull and left to dry. A silicon elastomer (Kwik-Cast) was then applied over the surface to prevent deterioration of skull transparency prior to photostimulation. Buprenorphine was used as a post-operative analgesic and the mice were allowed to recover over 5-7 days following surgery with free access to water.

For silicon probe recording, a small craniotomy of about 600 µm in diameter was made for implantation of a ground screw. The skull was thinned using a dental drill until the remaining bone could be carefully removed with a tungsten needle and forceps. Following this, one or more craniotomies of about 1 mm in diameter were made over the sites of interest for silicon probe recording. Craniotomies were protected with a layer of silicon elastomer (Kwik-Cast) on top. Additional craniotomies were usually made in new locations after finishing recordings in previous ones.

1.5.3   Task Control

Task control was implemented with an Arduino-based system (Teensy 3.2 and Teensyduino), including the generation of audio (Teensy Audio Shield). Custom MATLAB-based software with a graphical user interface was developed to log task events and change task parameters. Touches between the tongue and the port were registered by a conductive lick detector (Svoboda lab, HHMI Janelia Research Campus), where the mouse acted as a mechanical switch that opened (no touch) or closed (with touch) the circuit. Any mechanical switch has electrical bouncing issues when a contact is weak and unstable. To handle bouncing during loose touches, we merged any contact signals with intervals less than 60 ms.

The auditory cue that signaled the beginning of each trial was a 0.1 s long, 65 dB SPL, 15 kHz pure tone. Touches that occurred during the auditory cue were not used to trigger port movement as they were likely due to impulsive licking rather than a reaction to the cue.

The lick port was motorized in the horizontal plane by two perpendicular linear stages (LSM050B-T4 and LSM025B-T4, Zaber Technologies), one for anterior and posterior movement and the other for left and right. A manual linear stage (MT1/M, Thorlabs) installed in the vertical direction controlled the height of the lick port. The motors were driven by a controller (X-MCB2, Zaber Technologies) which was in turn commanded by the Teensy board via serial interface communication. Although the linear stages were set up in cartesian coordinates, we specified the movement of the port using a polar coordinate system. For a chosen origin of the polar coordinates, the seven port positions were arranged in an arc symmetrical to the midline with equal spacing (in arc length) between adjacent positions (Figure 2.1a). A movement of the lick port was triggered by the onset of a touch during sequence performance. A second port movement could not be triggered within 80 ms, which prevented mice from driving a sequence by constantly holding the tongue on the port (although we never observed such behavior). When a movement was triggered, the port first accelerated (477 or 715 mm/s^2) until the maximal velocity (39.3 mm/s) was reached, then maintained the maximal velocity, and decelerated until it stopped at the end position. The acceleration and deceleration phases were always symmetrical, such that the maximal velocity might not be reached if the distance of travel was short. The movement was typically in a straight line. For 4 of the 9 mice, when the two positions were not adjacent (e.g. at backtracking and the following transition), the port would move in an outward half circle whose diameter was the linear distance separating the two positions. This arc motion minimized the chance of mice occasionally catching the port prematurely before the port stopped. Nevertheless, catching the port prematurely did not trigger the next transition in a sequence because, in this case, the port movement could only be triggered again after 200 ms from the start of backtracking (and 300 ms after the following touch). As a result, mice always needed to touch the port at the fully backtracked position in order to continue progress in a sequence.

Mice performed the task in darkness with no visual cues about the position of the port. To prevent mice from using sounds emitted by the motor to guide their behavior, we played

a Sensor for horizontal displacement
Lick tube (cross-section)
Sensor for vertical displacement
Anchor

b $R^2 = 0.9999$
Force (mN)
Voltage (V)
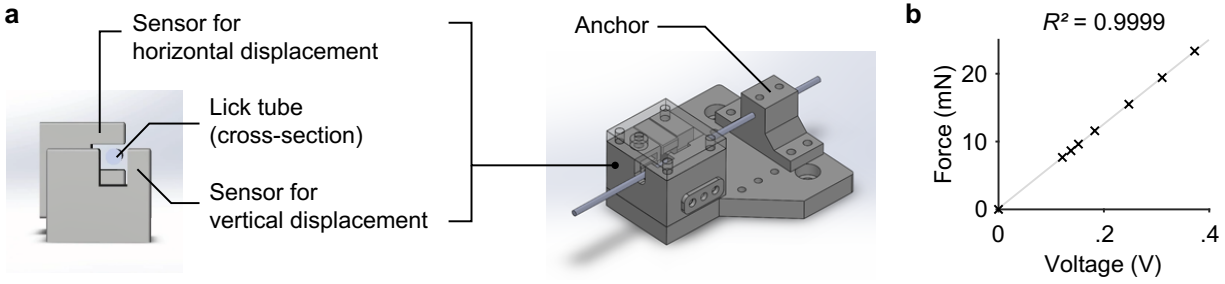
Figure 1.5: Two-axis optical force sensors

(a) CAD images of the sensor core (left) and the assembly (right) with a lick tube.

(b) Linear relationship between the applied force and the sensor output voltage.

two types of noise throughout a session. The first was a constant white noise (cutoff at 40 kHz; 80 dB SPL) and the second was a random playback (with 150-300 ms interval) of previously recorded motor sounds during 12 different transitions.

1.5.4   Two-axis Optical Force Sensors

A stainless steel lick tube was fixed on one end to form a cantilever. Mice licked the other free end, producing a small displacement ($< \sim 0.1$ mm at the tip for 5 mN) of the tube. Two photointerrupters (GP1S094HCZ0F, Sharp) placed along the tube (Figure 1.5a) were used to convert the vertical and horizontal components of displacement into voltage signals. Specifically, the cantilever normally blocked about half of the light passing through, outputting a voltage value in the middle of the measurement range. Pushing the tip down caused the cantilever to block more light at the vertical sensor and thereby decreased the output voltage; conversely, less force applied at the tip resulted in increased voltage. For the horizontal sensor, pushing the tube to the left or right decreased or increased the voltage output, respectively. Output was amplified by an op-amp then recorded via an RHD2000 Recording System (Intan Technologies).

By design (the circuit diagram and the displacement-response curve are available in the GP1S094HCZ0F datasheet), the force applied at the tip of the lick tube and the displacement

sensor's output voltage follow a near linear relationship within a range of forces. To find this range, we measured the voltages (relative to baseline) with different weights added to the tip. Excellent linearity ($R^2 = 0.9999$) was achieved up to >20 mN (Figure 1.5b). In contrast, the maximal force of a lick was on average about 4 mN (Figure 2.4a).

The motorization of the lick tube introduced mechanical noise to the force signals. The spectral components of these noises were mainly at 300 Hz and its higher harmonics, presumably due to the resonance frequency of the tube, whereas the force signal induced by licking occupied much lower frequencies. Therefore, we low-pass (at 100 Hz) filtered the original signal (sampled at 30 kHz/s) to remove the motor noise. Additional interference came from the 850 nm illumination light used for high-speed video, which leaked into the optical sensors (mainly in early experiments with 2 mice) and caused slow fluctuations in the baseline over seconds. To mitigate this slow drift, we used a baseline estimated separately for each individual lick as follows. We first masked out the parts of the signal when the tongue was touching the port, then linearly interpolated to fill in these masked out lick portions using the neighboring (i.e. no touch) values. These interpolated time series served as the baseline for each lick. Since the lick force was only a function of voltage change compared to baseline (rather than the absolute voltage), the above procedure would at most negligibly affect the force estimation. Due to the dependency of this procedure on complete touch detection, we excluded 8 sessions from behavioral quantifications in Figs. 1 and 2 where only touch onsets were correctly registered.

1.5.5 High-speed Videography and Tongue Tracking

High-speed video (400 Hz, 0.6 ms exposure time, 32 µm/pixel, 800 pixels × 320 pixels) providing side- and bottom-views of the mouth region was acquired using a 0.25X telecentric lens (55-349, Edmund Optics), a PhotonFocus DR1-D1312-200-G2-8 camera, and Streampix 7 software (Norpix). Illumination was via an 850 nm LED (LED850-66-60, Roithner Laser) passed through a condenser lens (Thorlabs).
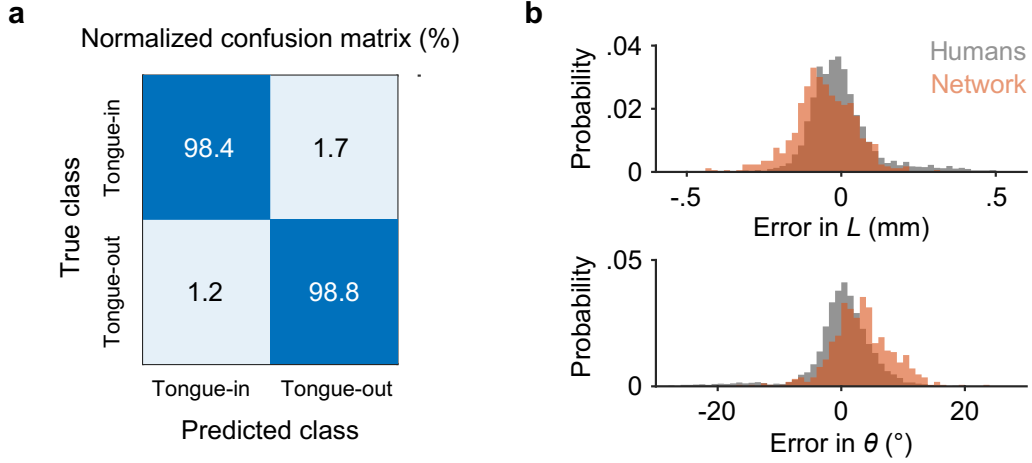
Figure 1.6: Performance of the deep neural networks for tongue tracking

(a) Confusion matrix showing the performance of the classification network. The numbers represent percentages within each (true) class (n = 1611 frames).

(b) Performance of the regression network. Top, the gray probability distribution shows how L from five human individuals varied from the mean L across the five. The red distribution shows how predicted L varied from the human mean. Bottom, similar quantification as the top but for θ. n = 573 frames.

Three deep convolutional neural networks were developed (MATLAB, Deep Learning Toolbox) to extract tongue kinematics and shape from these videos. The first network classified each frame as "tongue-out", if a tongue was present, or "tongue-in" otherwise. This network was based on a pretrained network, ResNet-50 (He et al., 2015), but the final layers were redefined to classify the two categories. A total of 37658 frames were manually labeled in which 1611 frames were set aside as testing data. Image augmentation was performed to expand the training dataset. A standard training scheme was used with a mini-batch size of 32 and a learning rate of 1×10^{-4} to 1×10^{-5}. The fully trained network achieved a high accuracy in classifying the validation data (Figure 1.6a).

The second network assigned a vector from the base to the tip of the tongue in each frame classified as "tongue-out". L and θ were derived from this vector (Fig. 1c). A total of 12095 frames were manually labeled in which 643 frames were used only for testing. The architecture and training parameters of this network are similar to those of the classification

network except that the final layers were redefined to output the x and y image coordinates of the base, tip and two bottom corners (not used in analysis) of the tongue with mean absolute error loss. The regression error of the fully trained network in testing data was 3.1 $\pm$ 5.4° for θ and 0.00 $\pm$ 0.13 mm for L (mean $\pm$ SD). This performance was comparable to human level (Extended Data Fig. 1b). Specifically, a subset of frames (separate from testing data) were labeled by each of the five human labelers. The variability in human judgement was quantified by the differences between L and θ from individual humans and the human mean for each frame. We also computed the differences between L and θ from the network and the human mean for each frame. The two distributions showed a comparable variability, although the network showed small biases (L: humans 0 $\pm$ 0.11 mm, network -0.05 $\pm$ 0.10 mm; θ: humans 0 $\pm$ 5.7° SD, network 3.3 $\pm$ 5.5° SD; mean $\pm$ SD).

In a subset of trials and in frames classified as "tongue-out", the third network, a VGG13-based SegNet (Badrinarayanan et al., 2015), extracted the shape of the tongue by semantic image segmentation, i.e. classifying each pixel as belonging to a tongue or not. Human labelers used a 10-vertex polygon to encompass the area of the tongue in a total of 3856 frames. The training parameters were similar to the other networks except for a mini-batch size of 8 and a learning rate of 1×10^{-3}.

1.5.6 Behavioral Training

Behavioural sessions occurred once per day during the dark phase and lasted for approximately an hour or until the mouse stopped performing, whichever came earlier. Mice would receive all of their water from these sessions, unless it was necessary to supply additional water to maintain a stable body weight. The amount of water consumed during behaviour was measured by subtracting the pre-session volume of water in the dispenser from the post-session volume. On days where their behaviour was not tested they received 1 ml of water. Mice were water restricted (1 ml/day) for at least 7 days prior to beginning training. Whiskers and hairs around the mouth were trimmed frequently to avoid contact with the

port.

The precise position of the implanted headpost varied across mice, so each mouse required an initial setup of the lick port's positions. The lick port moved in an arc with respect to a chosen origin (see Task Control). The origin was initially set at the midline of the animal and 2 mm posterior from the posterior face of the upper incisors. If there was any yaw of the head, the whole arc was rotationally shifted accordingly. The lick port's z-axis was manually adjusted until the lick port was approximately 1 mm below the interface between upper and lower lips when the mouth was closed.

In initial training sessions, the distance between the leftmost (L3) and the rightmost (R3) lick port position was reduced, the radius of the arc was shortened, and the water reward was larger. As mice learned the task, both the L3 to R3 distance and the radius of the arc were gradually increased over a few days of training (Figure 2.6b). The difficulty of the task was increased whenever the mouse showed improvements in performing the task at a given set of parameters. The difficulty remained constant in two conditions: either when the maximum set of parameters had been met (a radius of 5 mm for males and 4.5 mm for females) or if the mouse appeared demotivated (typically indicated by a significant decrease in the number of trials and licks). During the initial training sessions, water was occasionally supplemented at other points during the sequence to encourage licking behaviour. The amount of water reward per trial was eventually lowered to $\sim$3 µL. For 3 of the 24 mice included in this study, we first trained them to lick in response to the auditory cue with the lick port staying at fixed positions. After mice responded consistently to the go cue, we shifted to the complete task with gradually increased difficulty. Although the 3 mice performed similarly to others when well trained, this procedure proved to be less efficient than beginning with the complete task.

Once a mouse had become adept at standard sequences, they were trained on the back-tracking sequences. The first 9 fully trained mice were used in backtracking related analyses; later mice used for other purposes were not always fully trained in backtracking. For 5 of the 9 mice, we first trained them with backtracking trials in only one direction and added

the other direction once they mastered the first. For 3 of the 9 mice, backtracking trials and standard trials were organized into separate blocks of 30 trials each. In developing this novel task, we tested subtle variations in the detailed organization of trial types, such as varying the percentage of backtracking trials in a block, or different forms of jumps in the port position. Two of these 3 mice continued to performed the block-based backtracking trials during recording sessions.

1.5.7 Trial Selection

The first trial and the last trial were always removed due to incomplete data acquisition. Trials in which mice did not finish the sequence before video recording stopped were excluded from the analyses that involved kinematic variables of tongue motion.

1.5.8 Behavioral Quantification

The duration of individual licks was variable. In order to average quantities within single licks, we first linearly interpolated each quantity using the same 30 time points spanning the lick duration (from the first to the last video frames of a tracked lick). L' was computed before interpolation. When the tongue was short, the regression network showed greater variability in determining θ and sometimes produced outliers. Thus, we only show θ when L is longer than 1 mm. In addition, any "lick" with a duration shorter than 10 ms was excluded.

The instantaneous lick rate was computed as the reciprocal of the inter-lick interval. The instantaneous sequence speed was defined as the reciprocal of the duration from the touch onset of a previous port position to the touch onset of the next.

Directly averaging trials pooled across animals assumes that different animals follow the same distribution, and thus underestimates potentially meaningful animal-to-animal variability. To incorporate this variability, we performed a hierarchical bootstrap procedure (Davison and Hinkley, 1997) when computing confidence intervals where noted. In each

34

iteration of this procedure, we first randomly sampled animals with replacement, then, from each of these resampled animals, sampled sessions with replacement, and then trials from each of the resampled sessions. The statistic of interest was then computed from each of these bootstrap replicates. In our optogenetic inhibition experiments, each animal only contributed one behavioral session for a given experimental condition. Therefore, the hierarchy only had two levels.

1.6   Disclosures

One manuscript (Xu et al., 2020) posted to a preprint website has been reformatted to comply with the requirements of this thesis.

Chapter 2

Sequence Licking Task

2.1 Introduction

Sequence behaviors are ubiquitous in our everyday life. Everything related to language (speaking, writing/typing) is performed in sequences. We look at the visual world with sequences of saccades (Hayhoe and Ballard, 2005). Your morning routine is largely a memorised sequence of sequences. However, if you found yourself in a hotel instead of your home the sequences can be adjusted or even reprogrammed flexibly and immediately. Perhaps it is exactly because of the ubiquity and our excellence of sequence behaviors that an average person tends to take the production of such behaviors for granted or is even not aware of their existence, until brain regions responsible for sequence control start to dysfunction (reviewed in Chapter 1.2.1).

From a reductionist point of view, it is reasonable to begin understanding complex behaviors by first studying the individual components (Clower, 1998). Indeed, most studies in motor control and sensorimotor transformation have been focusing on simple (not necessarily in terms of the underlying mechanisms) stimulus-response associations (Crochet et al., 2019). However, many sequence behaviors have their unique features (reviewed in Chapter 1.1.1), such as an intrinsically hierarchical organization, the need for inference and flexible

control as sequences unfold. Even behaviors that can in principle be achieved by a chain of simple associations were shown to be controlled by the brain in a framework of sequences (Tanji, 2001).

Past research on the neural basis of sensorimotor sequences has employed a variety of model systems, though mainly with manual or saccadic sequences in non-human primates (Desrochers et al., 2016). For a detailed understanding of the circuit mechanisms, a mouse model with similar richness in behavior would be useful given its anatomical tractability and genetic accessibility. More importantly, most of the previous studies have focused separately on the memory-guided motor sequence or a chain of sensorimotor associations (e.g. Shima et al. (1996); Isoda and Tanji (2002); Berdyyeva and Olson (2010)). Therefore, the field would benefit from a paradigm where ongoing sequences can be simultaneously guided by both memory and sensory cues through a process of inference.

This section serves to provide a background for the novel behavioral task. The above is a typical introduction which follows the vein of past research. However, it is a hindsight that fails to reflect the actual thinking leading to the present task design. I believe that showing the trajectory of development and discovery can be helpful to understand why the task is the way it is, and be inspiring for future exploration. Previously, in a separate project, I trained mice to lick or not to lick based on complex cues. As a common "trick" in training, the distance between the lick port and the mouth was occasionally adjusted to balance the motivation and impulsiveness of mice licking. It appeared to me that when the port distance changed, the lick patterns gradually adapted as a result. Inspired by a study showing that mice could adapt to a force field when operating a manipulandum (Mathis et al., 2017), Yuxi Chen and I, out of curiosity, decided to formally test whether the strength of licks adapts to a change in target distance. Despite great efforts, the preliminary results did not seem interpretable and the behavioral design appeared very limited. Nevertheless, I developed the optical force sensors and pioneered video analysis during this period. Moreover, the observations certainly drew our attention to the nuances of the tongue sensorimotor control.

One day, when reading a theoretical paper about hippocampal spatial maps (Stachenfeld et al., 2017), I realized a computational resemblance between a rat running along a track and a sequential motion of an external object, such as a lick port. Both involve linking one sensorimotor state to another; both can have branching points, spatially versus procedurally. At this point, we had the general idea about mice sequentially licking across a number of positions, receiving a reward, and licking back - just like a rat navigating back and forth on a linear track. We quickly realized that asking mice to actively trigger, rather than follow, the movement of the lick port was 1) more analogous to common real-world interactions, 2) would explicitly involve a sensorimotor loop, and 3) was also more conducive to training. This basic form would later become the "standard" sequences. Inspired by the spatial teleportation in a video game for human subjects to play in an fMRI scanner (Deuker et al., 2016), we designed the "backtracking" sequences that "teleports" the lick port from one position back to a previous non-adjacent position. It requires new associations to be built and executed in a context-dependent manner.

2.2 Methods

The use of mice, surgical procedures, behavioral task control, instrumentation, behavioral training procedures, and common analysis methods are described in Chapter 1.5.

2.2.1 Hearing Loss

Hearing loss experiments were performed to exclude the possibility that mice used sounds produced by the motors to localize the motion of the lick port during sequence performance. To induce temporary hearing loss ($\sim$27.5 dB attenuation) (Mowery et al., 2015), we inserted two earplugs made of malleable putty (BlueStik Adhesive Putty, DAP Products Inc.) into the ear canal openings bilaterally under microscopic guidance. Earplugs were shaped like balls and then formed appropriately to cover the unique curvature of each ear canal. When

necessary, the positioning of the earplugs was readjusted, or larger balls were inserted. Five well trained mice performed one "earplug" session and one control session. Mice did not have experience with earplugs prior to the earplug session. In earplug sessions, mice were first anesthetized under isoflurane to implant earplugs (taking 11-12.5 mins), then were put back to the homecage to recover from anesthesia (taking 10-11.5 mins), and performed the task after recovery. In control sessions, mice were anesthetized for the same duration (and to remove earplugs if necessary), and allowed to recover for the same duration before performing the task.

2.2.2 Odor Masking

Odor masking experiments were performed to exclude the possibility that mice used potential odors emanating from the lick port to localize its position during sequence performance. A fresh air outlet (1.59 mm in diameter) was placed in front of the mouse and aimed at the nose from ~2 cm away with ~45° downward angle. We checked the coverage of air flow (2 LPM) by testing whether a water droplet (~3 L) would vigorously wobble in the flow at various locations, and confirmed that both the nose and all the seven port positions were covered. Prior to the test session, head-fixed mice were habituated to occasional air flows when they were not performing sequences. In the test session, the air flow was turned off first and turned on continuously after the 100th trial (in four mice) until the end of the session, or turned on first and turned off after the 100th trial (in two mice). The air-off period served as the control condition for air-on period.

2.2.3 Learning Curves

Learning curves included data from all trials except for the first and the last trial of each session. Values in the learning curves were averaged in bins of 100 trials, with 50% overlap of consecutive bins.

2.3 Results

2.3.1 Sequence Licking Task

We trained head-fixed mice to perform a task in which they used sequences of directed licks to advance a motorized port through 7 consecutive positions, either from left to right or right to left, after an auditory cue (15 kHz, 0.1 s) that signaled the start of a trial (Figure 2.1a). Each transition from one position to the next was driven in a closed-loop manner by
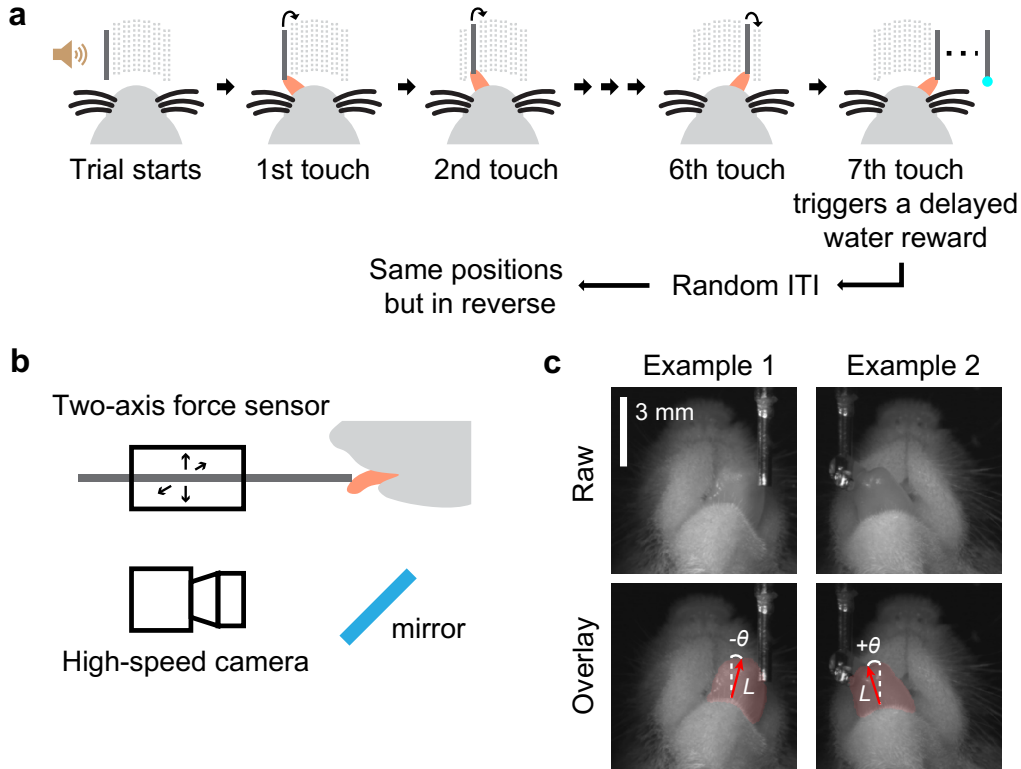
Figure 2.1: Behavioral setup

(a) Schematic of the sequence licking task.

(b) Schematic of the two-axis optical force sensors and high-speed videography in relation to a head-fixed mouse.

(c) Zoomed in high-speed video images showing the bottom view of a mouse licking the port. Overlaid vectors in red are outputs from the regression DNN, which point from the base to the tip of the tongue. Tongue length (L) is defined by the vector length. Tongue angle (θ) is the rotation of the vector from midline (positive being to the mouse's right). Light red shading depicts shape of the tongue based on output from the segmentation DNN.

a single lick touching the port. Thus, if a lick missed the port, the port would remain at the same position until the tongue eventually made contact. The port was no longer movable after the mouse had finished the 7 positions and a water droplet was delivered as a reward after a short delay (0.25 s, or 0.5 s in two mice). The next trial would then start after a random inter-trial interval (ITI) with a mean duration of 6 s, and the sequence would go in the opposite direction.

We used high-speed (400 Hz) video to capture tongue motions during sequence performance (Figure 2.1b) and developed an analysis pipeline based on deep artificial neural networks (DNN) to extract horizontal tongue angle (θ) and length (L) from each video
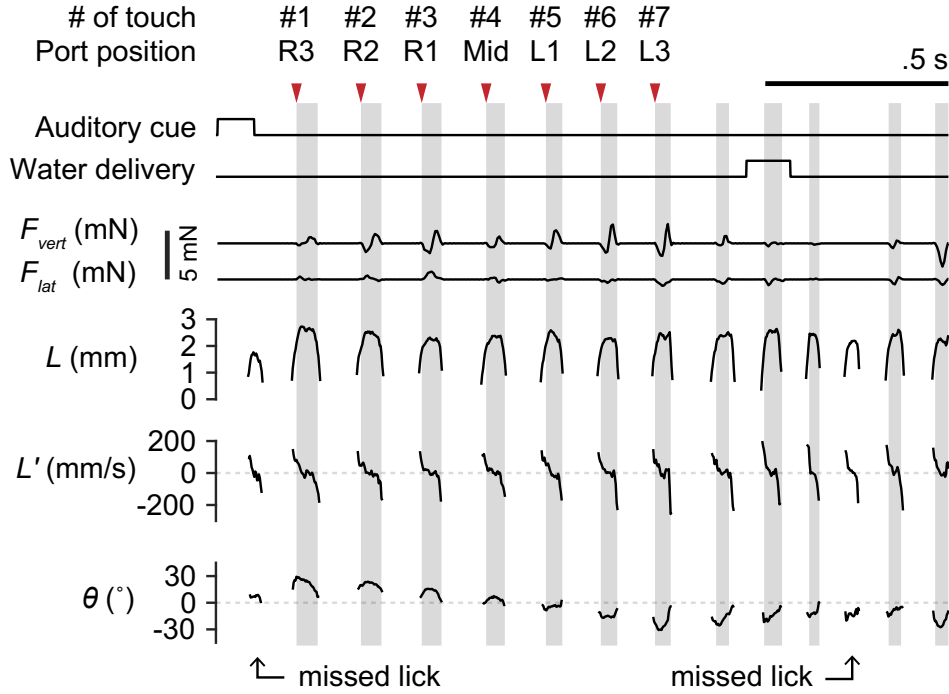
Figure 2.2: Behavioral measurements of an example trial

Time series of task events and behavioral variables during the first 2 s of an example trial. Variables recorded from the force sensors include the vertical lick force (F_{vert}, positive acts to lift the port up) and the lateral lick force (F_{lat}, positive acts to push the port to the right). Kinematic variables including L, its rate of change (L') and θ were derived from high-speed video. Periods of tongue-port contact are shaded in gray and are numbered ($\#$) sequentially. R3, R2, R1, Mid, L1, L2 and L3 indicate the 7 port positions from the rightmost to the leftmost. Arrowheads indicate the touch onsets for which a port movement (or the delayed water delivery at the last position) was triggered.

frame (Methods; Figure 2.1c). To quantify the sensory feedback experienced by the tongue when touching the port, we developed a set of sensors that measured the vertical and lateral component of instantaneous contact force (F_{vert} and F_{lat}) (Figure 2.1b and 1.5). Electrical contact detection was used to determine the precise onset and offset times of touch. An example trial shows all the measurements aligned in time (Figure 2.2).

Mice modulated each lick differently to reach different target locations (Figure 2.3). Specifically, the modulation was mainly in θ whereas the patterns of L and its rate of change (L') remained similar across targets (Figure 2.4a). When focusing on a single pattern at most lateral positions, we saw the tongue shooting out and quickly, but only briefly, reaching maximal deviation from midline ($|\theta - 0°|_{max}$ or simply $|\theta|_{max}$) (Figure 2.4b). As a result, the onset of touch mostly occurred around $|\theta|_{max}$. Interestingly, the modulation of L did not match that of $|\theta|$, suggesting a potential dissociation of control. Later, when analyzing licks which may or may not have contact, we use θ_{shoot}, defined as the θ when L reaches 0.84 maximal L (L_{max}), to succinctly depict the lick angle (Figure 2.4b). In addition, we will use capital Θ to represent unified tongue angles where the sign in right to left sequences is flipped so that data with both sequence directions can be pooled together.

Mice performed the task in darkness such that there were no visual cues to guide the licks.
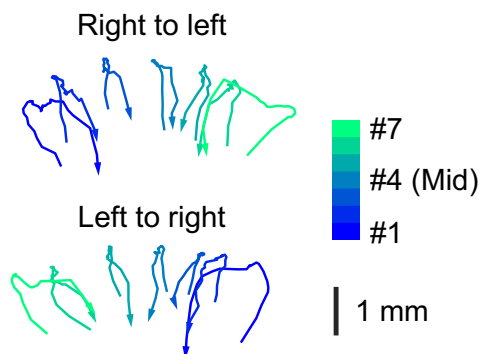
Figure 2.3: Example trajectories of two licking sequences

Two example trials showing the trajectories of the tongue tip when a mouse sequentially reached the 7 port positions, for both sequence directions. Arrows indicate the direction of time within each trajectory.
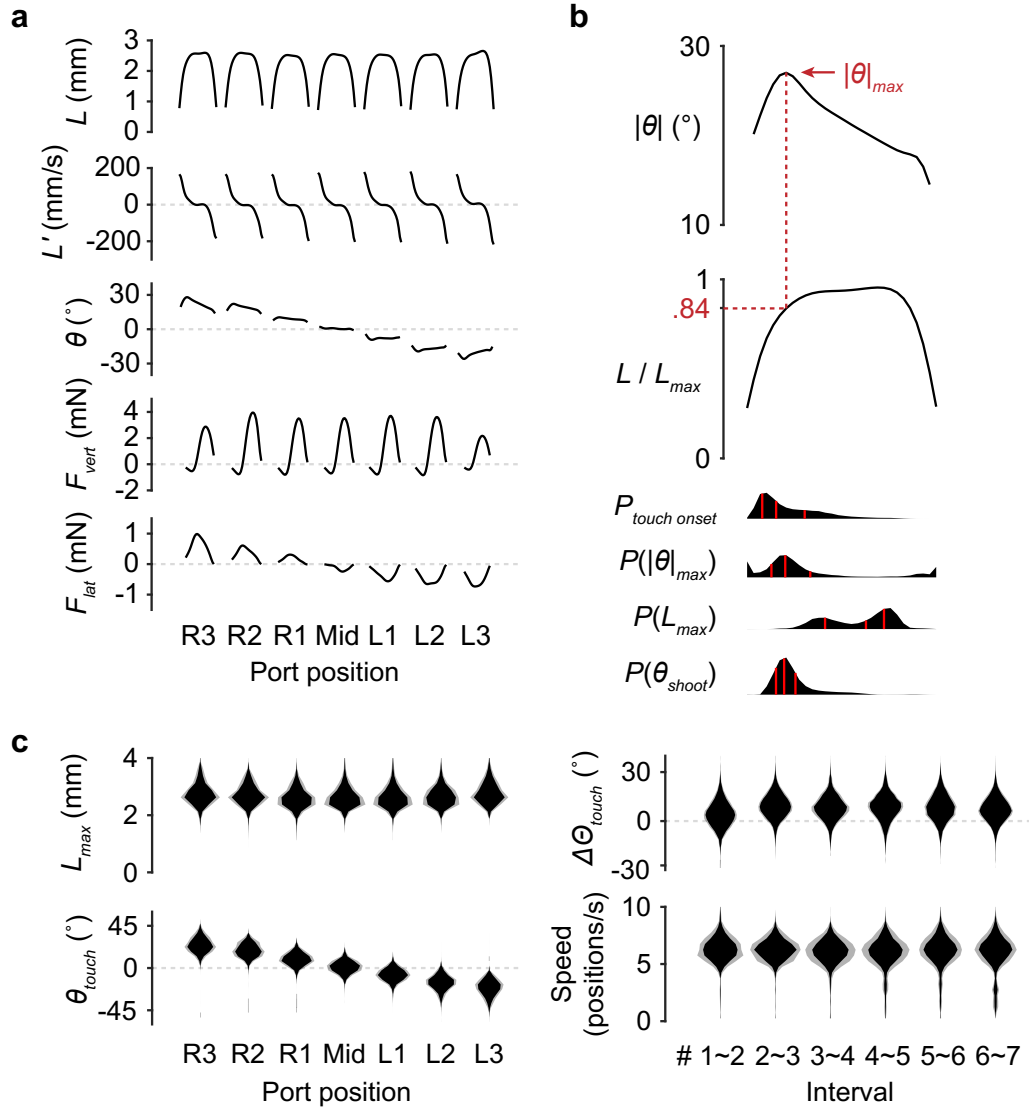
Figure 2.4: Quantification of sequence licking patterns

(a) Patterns of kinematics and forces of single licks at each port position (n = 22003 trials; mean ± 95% bootstrap confidence interval). The duration of individual licks was normalized.

(b) Top, the pattern of angle deviation from midline ($|\theta - 0°|$ or simply $|\theta|$) of single licks at R3 and L3. The vertical line indicates maximum $|\theta|$ ($|\theta|_{max}$). Middle, tongue length (L) expressed as a fraction of its maximum (L_{max}). The horizontal line indicates, on average, the fraction where $|\theta|_{max}$ occured. Bottom, time aligned probability distributions showing when touch onset, $|\theta|_{max}$, L_{max} or θ_{shoot} occured. Red lines mark quartiles. n = 22003 trials. Lick patterns show mean ± 95% bootstrap confidence interval.

(c) Top, probability distributions of Lmax and touch for licks at each port position. Bottom, probability distributions of the change in Θ_{touch} ($\Delta\Theta_{touch}$) and instantaneous sequence speed (Methods) for each interval separating port positions. Distributions show mean ± SD across n = 15 mice.

43

We played both white noise and pre-recorded mechanical noise of port transitions as masking sounds to prevent mice from using auditory cues (Methods). In addition, temporarily induced hearing loss (Methods) rendered mice unable to respond to the auditory cues but did not affect sequence performance, except in one of five mice, compared with control (Figure 2.5a). To test whether or not mice used odor emanating from the port for localization, we masked
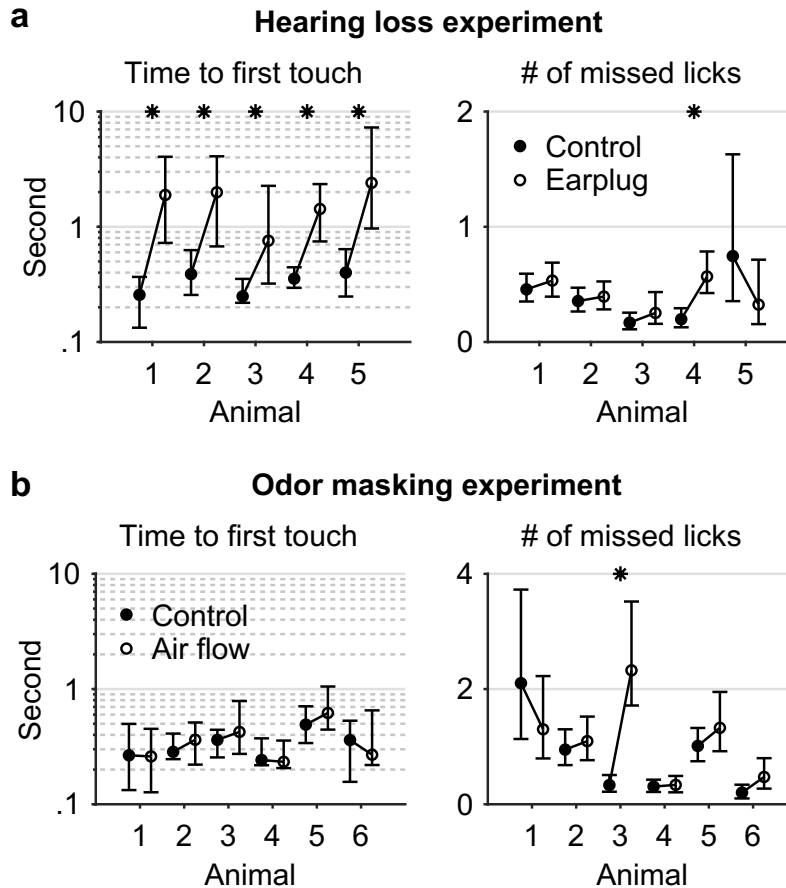
Figure 2.5: Sequence performance does not depend on audition or olfaction

(a) Time to first touch (left; median $\pm$ interquartile range) and number of missed licks during sequence performance (right; mean $\pm$ 95% bootstrap confidence interval) in control versus hearing loss (earplug) conditions. $*$ p < 0.01, not significant otherwise, one-tailed KS-test, Bonferroni correction for multiple comparison.

(b) Time to first touch (left; median $\pm$ interquartile range) and number of missed licks during sequence performance (right; mean $\pm$ 95% bootstrap confidence interval) in control versus odor masking (air flow) conditions. $*$ p < 0.05, not significant otherwise, one-tailed KS-test, Bonferroni correction for multiple comparison.

out any potential odor by applying a constant flow (2 LPM) of fresh air covering the nose and port region (Methods), and found no significant drop in sequence performance, except in one of six mice, compared with the condition without air flow (Figure 2.5b).

Mice typically obtained proficiency in standard sequences after a total of ∼1500 trials of training (Methods; Figure 2.6). In addition to stereotypic licking kinematics, expert mice showed remarkable speed of sequence execution, with the 7 positions completed in about a second (Figure 2.6a).
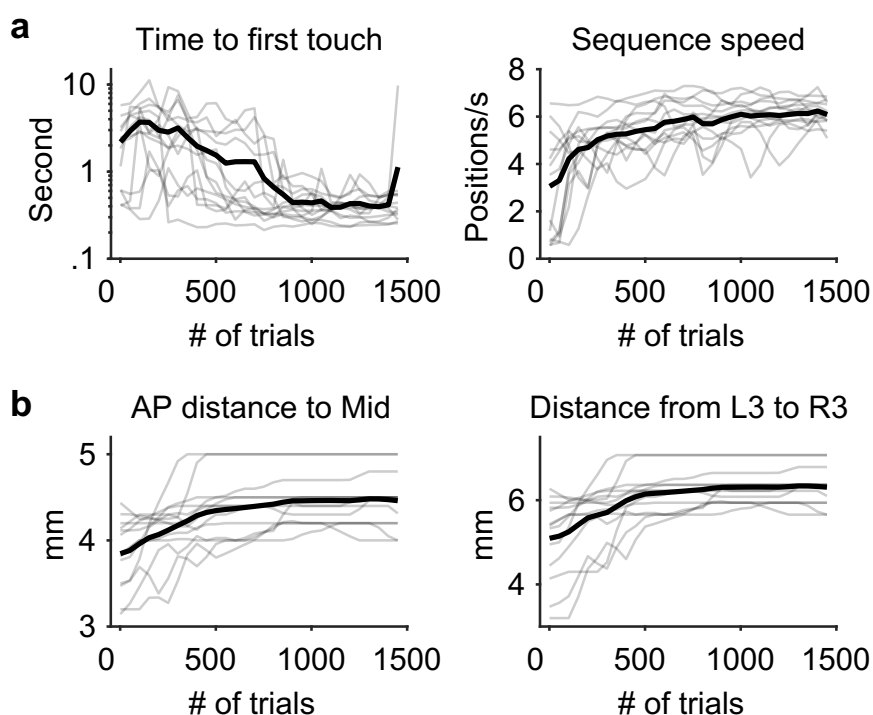
Figure 2.6: Learning curves for standard sequences

(a) Learning curves for 13 individual mice (gray) and the mean (black) showing the reduced sequence initiation time (left) in response to the auditory cue and the increased sequence speed (right).

(b) Gradual increase in task difficulty (Methods) accompanying the improved performance shown in (**a**).

2.3.2 Flexible Execution of Sequences with Backtracking

To determine if the sequence generation after training was strictly "ballistic" or was capable of flexible reconfiguration based on sensory feedback, we next varied the task by introducing unexpected port transitions after mice had learned the standard sequences (Figure 2.7a). Specifically, in a randomly interleaved subset (1/3 or 1/4) of trials, when a mouse licked at the middle position, the port would backtrack two steps rather than continue to the anticipated position (Figure 2.7b). Mice previously trained only with standard sequences learned to detect the change of port transition, lick to the new position and finish the rest of the sequence. However, mice showed mixed learning curves (Figure 2.8). About 3 mice were more biased to previously learned standard sequences and tended to miss the port
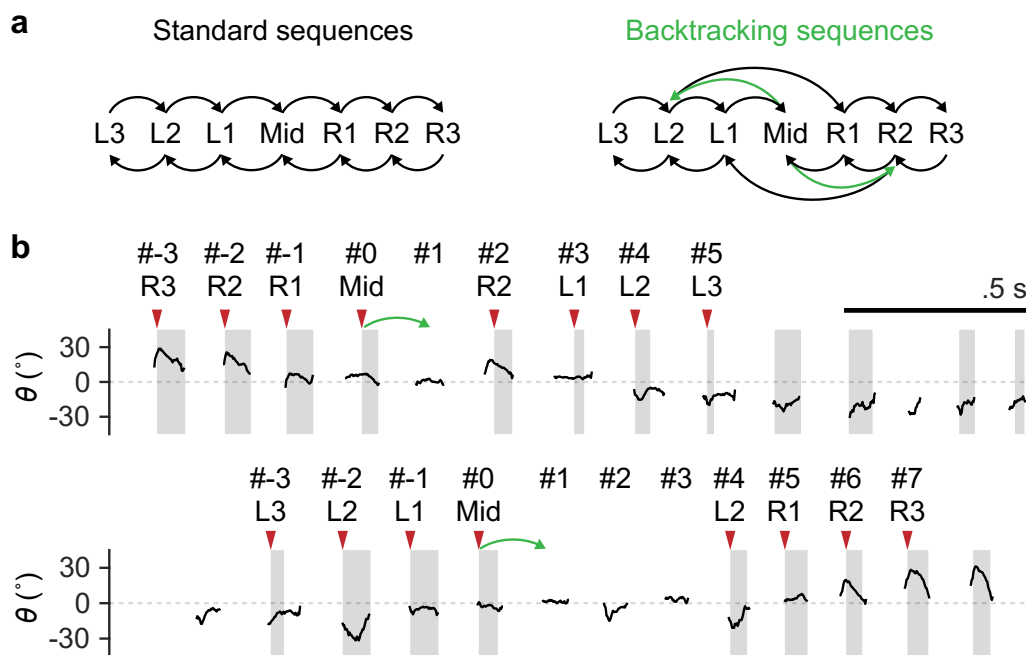
Figure 2.7: From standard to backtracking sequences

(a) Transition diagrams depicting standard sequences and those with backtracking (green).

(b) Example trials where the port backtracked (green arrows) when a mouse touched Mid. Licks including both touches and misses are indexed with respect to the lick at Mid. Top, the mouse missed once before it successfully relocated the port and finished the rest of the sequence. Bottom, after an initial miss and a lick back, the mouse made a second attempt and caught the port.
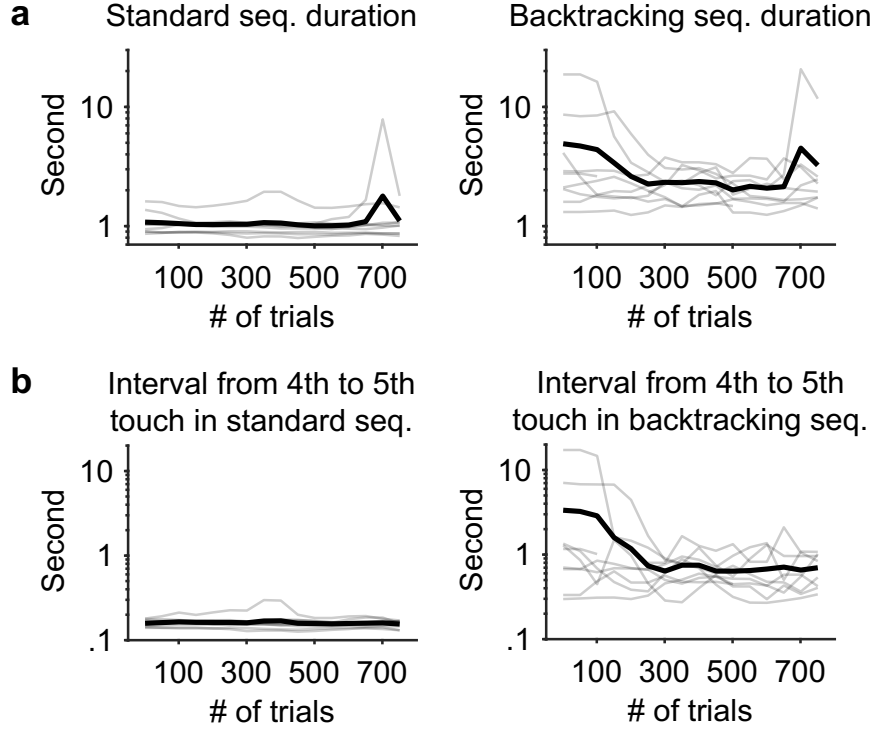
Figure 2.8: Learning curves for backtracking sequences

(a) Learning curves for 9 individual mice (gray) and the mean (black) showing the duration of time spent to perform standard (left) and backtracking (right) sequences.

(b) Similar to (**a**) but limited to the interval following the 4th (middle) lick in standard (left) or backtracking (right) sequences.

many times before relocating the lick port through exploration. The other 6 mice more readily made changes. On average, well-trained mice took 1 to 2 missed licks before quickly relocating the port (Figure 2.9a).

The lick patterns of L and L' appeared identical whether a mouse was performing a standard sequence or one with backtracking (Figure 2.9b). In contrast, Θ was strongly modulated toward backtracked directions starting at the second lick after backtracking occurred. A slight change in Θ was also present in the first lick after backtracking, presumably due to the mouse sensing the onset of port motion on its tongue toward the end of the prior lick. To our surprise, mice did not pause the ongoing sequence after detecting backtracking (Figure 2.9c). Instead, they showed the same, if not higher, lick rate when making adjustments in
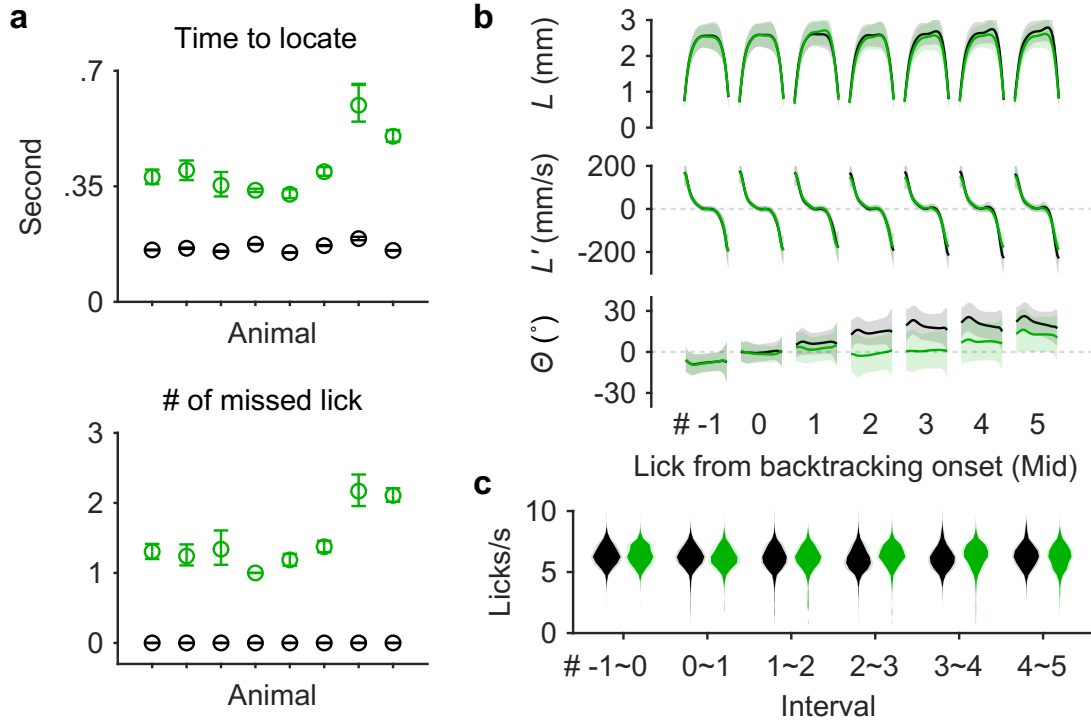
Figure 2.9: Flexible execution of sequences with backtracking

(a) Top, time to locate the port at its next position during the 4th interval, for standard sequences (black) or for sequences when the port backtracked (green). Bottom, the number of missed licks during the 4th interval. Mean ± 95% bootstrap confidence interval.

(b) L, L' and θ patterns for seven consecutive licks aligned at the 4th (Mid) touch (#0). Licks in standard sequences (n = 7365 trials) are shown in black, those in backtracking sequences (n = 2674 trials) are in green. Mean ± SD.

(c) Probability distributions of instantaneous lick rate for each interval separating consecutive pairs of the seven licks during standard (black) or backtracking (green) sequences (n = 8 mice; mean ± SD).

single lick bouts. This suggests that flexible control of sequence execution can be independent from the pattern generators responsible for generating basic rhythmic licking patterns (McElvain et al., 2018).

Together, these observations demonstrate that head-fixed mice can learn to perform complex and flexible licking sequences guided by sensory feedback.

2.4 Discussion

Rodents are highly tactile species and use rich orofacial behaviors to interact with the environment (McElvain et al., 2018). We found that mice exhibited a remarkable ability to learn complex licking sequences with behavioral features beyond simple stimulus-response associations.

The licking sequences have a hierarchical structure and require guidance from memory. Individual directed licks are organized in short sequences with two directions, which in turn are strung into a high-level sequence where the two directions alternate. Efficient sequence execution requires mice to know what is about to come by keeping a memory of the past. For example, to start a sequence promptly, mice reacted to the same auditory cue by initiating licks to different sides depending on which direction the sequence would go (also see Results in Chapter 5). In the memory-guided Pull-Push-Turn task (Tanji and Shima (1994); reviewed in Chapter 1.1.2), monkeys returned hand to the same standby position between movements. As a result, the end state of a previous movement could not be used to guide the next. In this task, mice retracted the tongue back to mouth with a demodulated tongue angle (Figure 2.4b) before protruding for the next lick angle. A memory of sequence context (e.g. sequence direction) may be required to determine which direction to go next. However, to test this more rigorously, one would need different sequences to contain the same subsequence of licks (instead of a single lick to be the same) to avoid any small differences in kinematics (Figure 2.3; middle licks).

The learned sequences are not immutable. Mice learned to use the absence or presence of tactile feedback to decide whether or not an ongoing sequence should be branched to a backtracking sequence. Modifying motor sequences by branching differs from classical cerebellum-dependent motor adaptation in at least two ways (also see Krakauer et al. (2019)). First, backtracking performance involved complete reversals of direction and reorganization of the motor sequence rather than fine adjustments to the standard sequence. Second, feedback-driven corrections could occur in a trial-by-trial and all-or-none fashion. In contrast,

cerebellum-dependent adaptations to perturbations usually show a gradual shift in movement kinematics over multiple trials and a washout effect after the perturbations are removed (Bastian, 2006; Shadmehr et al., 2010). However, this is not to suggest that cerebellum is not involved in other aspects of sequence control. It has been shown that the motor cortex, motor thalamus, and cerebellum form robust loops responsible for the preparation and control of directed licks (Gao et al., 2018; Chabrol et al., 2019).

Although the present study mainly used four different sequences, one can imagine a practically unlimited number of alternative sequences. Furthermore, different sequences do not have to be randomly interleaved or occur with specific probabilities. For example, to investigate how mice would behave with different prior (or expected) sequences, one can put the same sequences in different blocks of trials and study the performance at block transitions. Another interesting question is whether mice can generalize sequence transitions in one direction to the other in a mirrored fashion like what has been reported in primates (Jiang et al., 2018). Because mirroring changes all the physical relationships from one movement to another, an ability to generalize would suggest that the sequence structure is represented independently from detailed licking kinematics. To test the possibility, one can first train mice to learn a novel sequence and then introduce its mirrored sequence. If mice generalize the abstract sequence structure, the second phase of learning will be much faster than the first.

Overall, the novel sequential licking task fulfills the need for a rodent model to study complex sensorimotor sequences. It demonstrates how skillful and flexible mice can use their tongue, in coordination with other orofacial apparatus, to accomplish goals. The versatile task design opens up new opportunities for future research.

Chapter 3

Closed-loop Optogenetic Inhibition Screen of Cortical Areas

3.1   Introduction

Identifying brain regions that produce the behavior of interest is the first step to understand its neural underpinning. The present sequence licking task is a novel behavior, the first in rodents involving complex sensorimotor sequence control. It is not clear which cortical regions, if any, are responsible for the sequence generation and what their roles are?

The performance of licking sequences relies on the coordinated motor control of tongue, jaw and other orofacial apparatus (Travers et al., 1997; McElvain et al., 2018), but also tactile and perhaps proprioceptive feedback from these structures (Miller, 2002). Electrical and optogenetic mapping studies have generally charted two motor subregions in rodent brains that are able to evoke tongue and jaw movements (Svoboda and Li, 2018). One is the anterolateral motor cortex (ALM) (Komiyama et al. 2010); another more lateral to it is thought to be the tongue/jaw subregion of the primary motor cortex (or M1TJ) (Mayrhofer et al., 2019). Regarding the sensory processing, physical stimulation to the tongue evokes major hemodynamic and intracellular calcium responses in the tongue/jaw subregion of

primary somatosensory cortex (S1TJ), and a secondary response in M1TJ, presumably due to prominent axonal projections found from S1TJ to M1TJ (Mayrhofer et al., 2019). Inhibiting ALM and M1TJ neurons can bias the production of directed licks (Li et al., 2015; Allen et al., 2017; Mayrhofer et al., 2019), but little is known about the contribution of S1TJ. It is also not clear how the ALM, M1TJ, S1TJ, and potentially other cortical regions, would function when mice perform more complicated sensorimotor sequences. Therefore, a systematic loss-of-function screen of all the candidate regions is needed.

In this sequence licking task, mice wait for an auditory cue to initiate a sequence, then roll out a specific motor program based on memory and sensory feedback, and eventually terminate the licking sequence after the water reward is consumed. A brain region can be involved in all or only a subset of the three temporal stages. To dissect these contributions, the loss-of-function manipulation must have high temporal resolution and be triggered at the right time points. Optogenetic inhibition can be turned on and off rapidly and can be directed to specific brain regions and circuit components. Guo and colleagues pioneered the "clear-skull" technique which renders intact skulls optically transparent (Guo et al., 2014). Light can be delivered to any region of the dorsal cortex without invasive or irreversible preparations, thus making high throughput functional screening during behavior possible. Additionally, in a self-paced task like ours, the timing of optoinhibition must be triggered by the behavior in a closed-loop manner.

In sum, closed-loop optoinhibition across cortical regions combined with detailed behavioral quantification enable us to systematically characterize the role that each cortical region plays in the control of sequence licking.

3.2   Methods

The use of mice, surgical procedures, behavioral task control, instrumentation, behavioral training procedures, and common analysis methods are described in Chapter 1.5. Common

procedures for silicon probe recording and spike sorting are describes in Chapter 4.2.1.

3.2.1 Photostimulation

Bilateral stimulation of the brain was achieved using a pair of optic fibers (0.39 NA, 400 μm core diameter) that were manually positioned above the clear skull prior to the beginning of each behavioural session. These optic fibers were coupled to 470 nm LEDs (M470F3, Thorlabs). The illumination power was externally controlled via WaveSurfer (http://wavesurfer.janelia.org). Each stimulation had a 2 s long 40 Hz sinusoidal waveform with 0.1 s linearly modulated ramp-down at the end. The peak powers in the main experiments were 16 mW and 8 mW. We used the previously reported 50% transmission efficiency of the clear-skull preparation (Guo et al., 2014) and report the estimated average power in the Results. There was a 10% chance of light delivery triggered at each of the following points in a sequence: cue onset, the fourth (middle) touch, or the first touch after water delivery. To ensure that the light from photostimulation did not affect the mouse's performance through vision, we set up a masking light with two blue LEDs directed at each of the mouse's eyes. Each flash of the masking light was 2 s long separated by random intervals of 5-10 s. This masking light was introduced several training sessions in advance of photostimulation to ensure the light no longer affected the behaviour of the mouse. In addition, the optic fibers were positioned to shine light from ∼5-10 mm above the animal's head on these days leading up to photostimulation.

In a subset of silicon probe recording sessions (related to Figure 3.2), we used an optic fiber (0.3 NA, 400 um core diameter) to simultaneously photoinhibit the same (within 1 mm) or a different cortical region (∼1.5 or ∼3 mm away) via a craniotomy. The tip of the fiber was kept ∼1 mm away from the brain surface. For testing the efficiency of photoinhibition, the same 2 s photostimulation was applied but only at the mid-sequence, with 7.5% probability for each of the four powers (1, 2, 4 and 8 mW). For each isolated unit, the photo-evoked spike rate was normalized to that obtained during the equivalent 2 s time window without

photostimulation. To avoid a floor effect, we also excluded units that on average fired less than one spike during the no stimulation windows. We classified units as putative pyramidal neurons if the width of the average spike waveform (defined as time from trough to peak) was greater than 0.5 ms, and as putative fast spiking interneurons if shorter than 0.4 ms or if units had more than twice the firing rate during 8 mW photostimulations than during periods of no stimulation.

3.2.2 Behavioral Quantification

The rate of licks, rate of touches, Θ_{shoot} and L_{max} as a function of time were computed using 0.2 s time bins. The summary quantifications (Figures 3.7 and 3.9) used data averaged within 1 s after the start of photoinhibition (or the equivalent time in no-inhibition trials). The shorter window helped to minimize the effects "bleeding over" from mid-sequence to initiation, and from consumption to mid-sequence. Although this was not an issue for the consumption period, we nevertheless used the 1 s window for consistency. The summary metric SD(Θ_{shoot}) was obtained by averaging the standard deviation of Θ_{shoot} in each 0.2 s time bin within the 1 s window. Other metrics were directly computed without binning.

3.3 Results

To determine which brain regions contributed to the performance of our novel sequence licking task, and at which points during sequence execution, we performed systematic closed-loop optogenetic silencing experiments.

We used the "clear-skull" preparation (Guo et al., 2014), a method that greatly improves the optical transparency of intact skull, to non-invasively photoactivate channelrhodopsin-expressing GABA-ergic neurons and thus indirectly inhibit nearby excitatory neurons (Figure 3.1a). In different experimental sessions, bilateral inhibition was centered at each of five regions: the ALM (Svoboda and Li, 2018) cortex, a somatomotor region centered at the
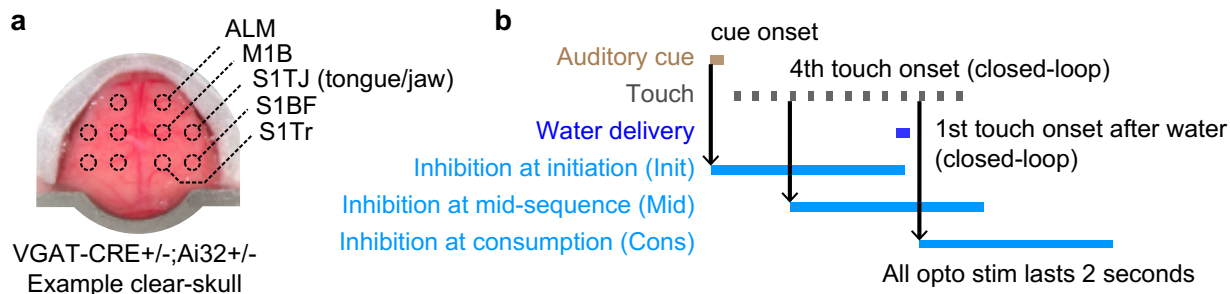
Figure 3.1: Experimental setup for optoinhibition screen

(a) Example clear-skull preparation overlaid with circles indicating the five bilateral target regions.

(b) Triggering scheme for photoinhibition at sequence initiation, mid-sequence and water consumption.

body-related primary motor cortex (M1B) (Ayling et al., 2009; Guo et al., 2015), the S1TJ cortex (Mayrhofer et al., 2019; Clemens et al., 2018), the macrovibrissae subregion (or barrel field) of the primary somatosensory cortex (S1BF), and the trunk subregion of the primary somatosensory cortex (S1Tr, also including a part of posterior parietal cortex). For each region, inhibition was triggered with equal probability (10%) at sequence initiation, at mid-sequence or at the start of water consumption (Figure 3.1b). Importantly, stimulation at mid-sequence and at consumption was triggered in closed-loop by the onset of the fourth (middle) touch during the sequence and by the onset of the first touch after water delivery, respectively. In the case of consumption, the trigger conditions ensured that the tongue had contacted water before inhibition started.

With the light powers we used (4 mW each hemisphere; Methods), light within 1 mm distance reduced mean spike rate of putative pyramidal cells (Figure 3.2a-c) by 91%, light at ∼1.5 mm away by 61%, and ∼3 mm away by 19% in behaving animals (Figure 3.2d). Interestingly, the mean spike rate of putative fast spiking (FS) neurons at ∼3 mm away was also reduced by 19%, rather than showing an increase due to photoactivation, suggesting that the decreased activity of both pyramidal and FS neurons was likely due to a reduction of cortical input. In contrast, light shined within 1 mm increased the mean spike rate of FS
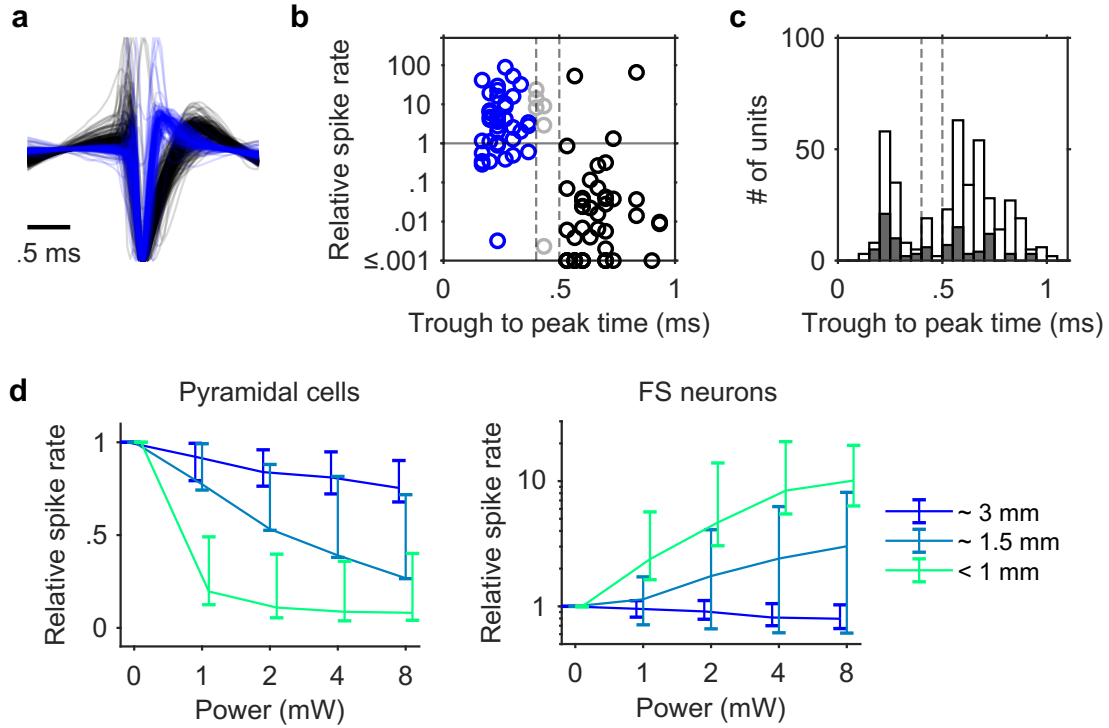
Figure 3.2: Efficiency of optoinhibition

(a) Average spike waveform of putative pyramidal cells (black; n = 224) and putative FS neurons (blue; n = 117), normalized to the amplitude of negative peaks.

(b) Relationship between spike widths (defined as the trough to peak time of average waveform) and changes in mean spike rate under opto illumination (4mW, within 1 mm) relative to baseline. Pyramidal cells (black; n = 42) and FS neurons (blue; n = 41) were classified by the two thresholds (dashed lines at 0.4 and 0.5 ms) with ambiguous units (gray; n = 6) in the middle.

(c) Distributions of spike widths from neurons in (**b**) (filled bars; n = 89) and from all neurons (empty bars; n = 414) including those where illuminations were not at recording sites. Classification thresholds are shown in dashed lines.

(d) Left, inhibition efficiency of putative pyramidal cells as a function of light power and distance away from the center of illumination (n = 224 units total). Right, similar to left but shows the excitation efficiency of putative FS neurons (n = 117 units total). Mean ± 95% hierarchical bootstrap confidence interval.

neurons by 739% and at ~1.5mm by 140%.

3.3.1 S1TJ is Required for Proper Targeting

Somatosensory inputs not only provide information about external objects but also enable the proprioceptive sensing of the body's position in space (Proske and Gandevia, 2012) for motor control (Franklin and Wolpert, 2011). Missing sensory feedback can render seemingly effortless manipulations surprisingly difficult even though motor capability per se is unchanged (Chesler et al., 2016). Inhibiting S1TJ randomized the targeting angle of licks, shown by the increased standard deviation (SD) in of individual licks (Figure 3.3a; shadings around traces) and the SD in Θ_{shoot} at given time points for individual mice (Figure 3.6d and 3.7d). Despite the increased variability in targeting, the ability to direct licks away from the midline (i.e. $|\Theta_{shoot} - 0°|$ or simply $|\Theta_{shoot}|$) was uncompromised (Figure 3.7c). Inhibiting S1TJ also did not shorten the length of licks (Figure 3.6b and 3.7b). Taken together, this suggests that S1TJ inhibition left intact the core motor capabilities required for tongue protrusions and licking, but corrupted these commands and their proper targeting, possibly due to missing sensory feedback.

In contrast, when inhibiting ALM (Figure 3.3b), mice not only had difficulty directing licks away from the midline (Figure 3.6c and 3.7c), but also showed decreased length of lick (Figure 3.6b and 3.7b). The variability of licks, however, did not change (Figure 3.6d and 3.7d). Inhibiting M1B (Figure 3.4a) only caused a moderate increase in the variability of lick angle and no decrease in lick length. Inhibiting S1Tr (Figure 3.4b) and S1BF (Figure 3.5 or 3.6) did not change any aspects of lick control (Figure 3.7).

3.3.2 ALM is Critical for Sequence Initiation

ALM has been shown to be important in motor preparation of directed single licks to obtain water reward (Komiyama et al., 2010; Svoboda and Li, 2018). Here, we found that inhibiting ALM at sequence initiation strongly suppressed the production of licking sequences (Figure
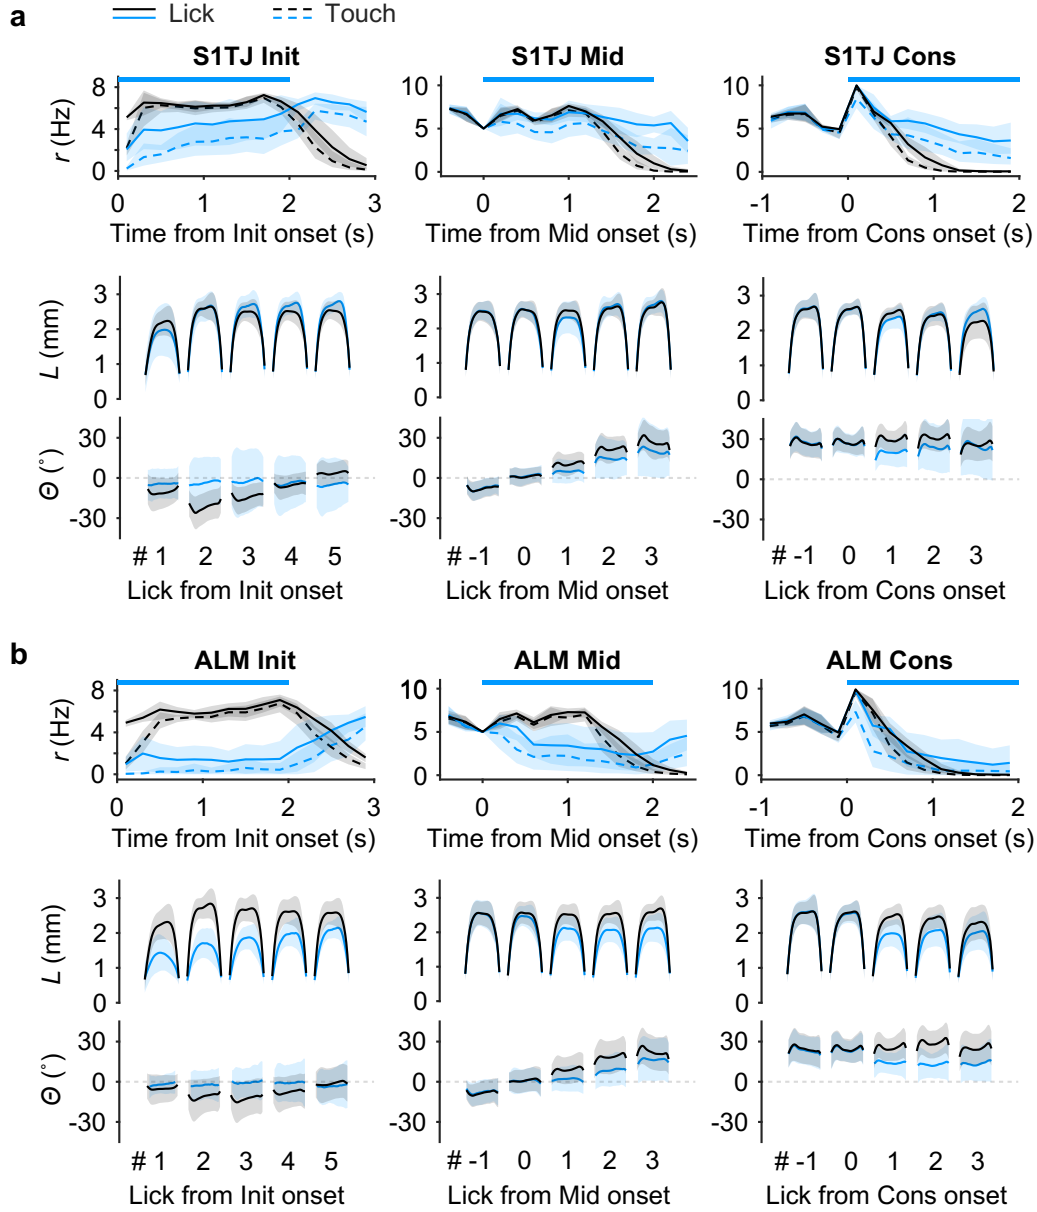
Figure 3.3: Effects of optoinhibition to S1TJ and ALM

Effects of S1TJ (**a**) and ALM (**b**) inhibition for each of the three periods. Top row, the rate of licks (solid lines) or touches (dashed lines) as a function of time for trials with (blue) or without (black) inhibition (mean ± 95% hierarchical bootstrap confidence interval). Horizontal bars indicate the 2 s long illumination. Time zero is the onset of inhibition, or the equivalent time in trials without inhibition. Middle and bottom rows, L and patterns (mean ± SD) for five consecutive licks in trials with (blue) or without (black) inhibition. Lick #0 is the lick that triggered inhibition, #1 the first lick after the trigger, #-1 the lick before, etc. Licks from inhibited trials are colored blue and those without inhibition are in black. n = 1542 trials for S1TJ from 5 sessions, one session per mouse. n = 1181 trials for ALM.
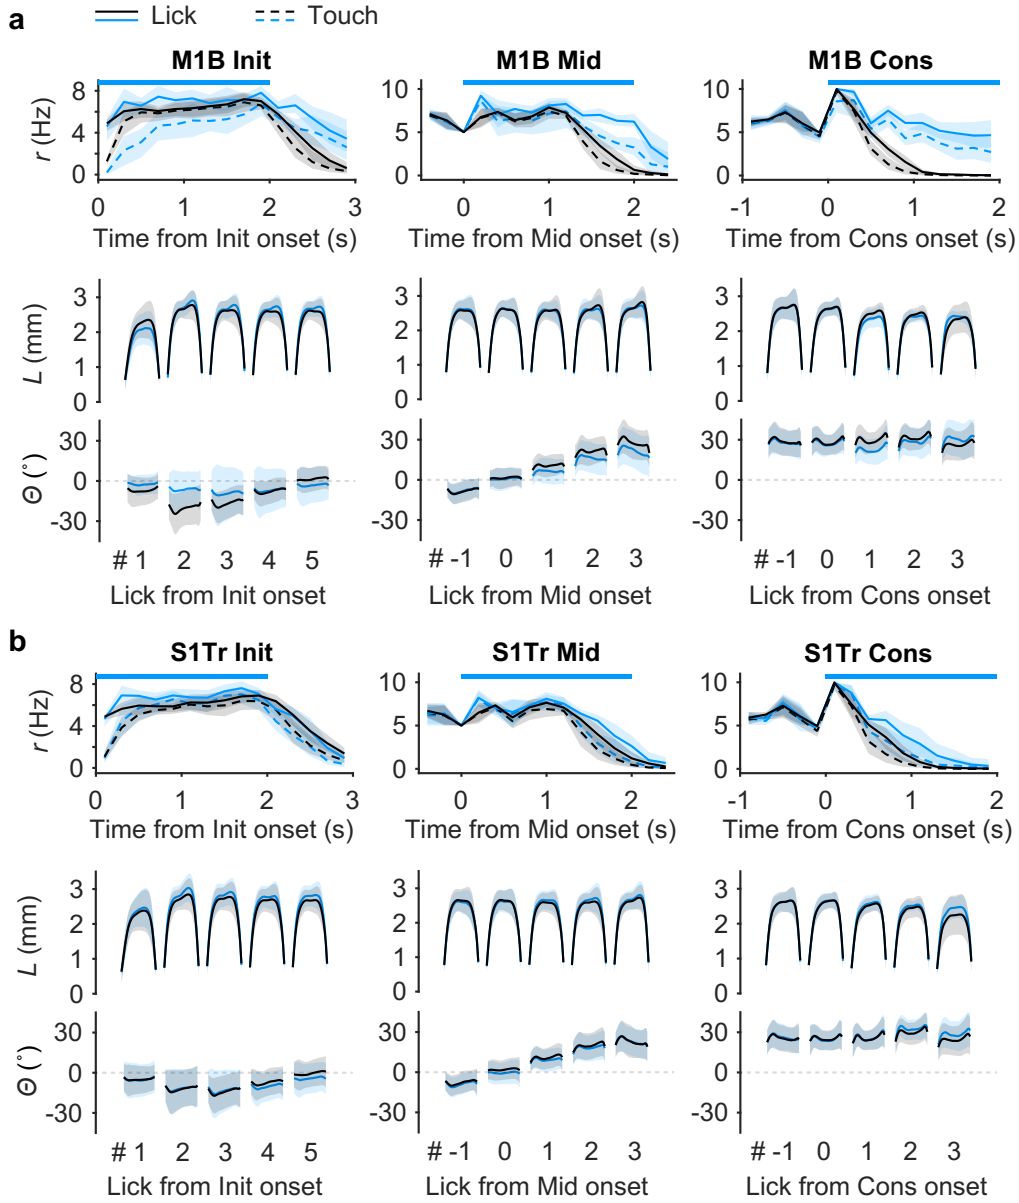
Figure 3.4: Effects of optoinhibition to M1B and S1Tr

Effects of M1B (**a**) and S1Tr (**b**) inhibition for each of the three periods. Top row, the rate of licks (solid lines) or touches (dashed lines) as a function of time for trials with (blue) or without (black) inhibition (mean $\pm$ 95% hierarchical bootstrap confidence interval). Horizontal bars indicate the 2 s long illumination. Time zero is the onset of inhibition, or the equivalent time in trials without inhibition. Middle and bottom rows, L and   patterns (mean $\pm$ SD) for five consecutive licks in trials with (blue) or without (black) inhibition. Lick #0 is the lick that triggered inhibition, #1 the first lick after the trigger, #-1 the lick before, etc. Licks from inhibited trials are colored blue and those without inhibition are in black. n = 1240 trials for M1B from 5 sessions, one session per mouse. n = 1269 trials for S1Tr.
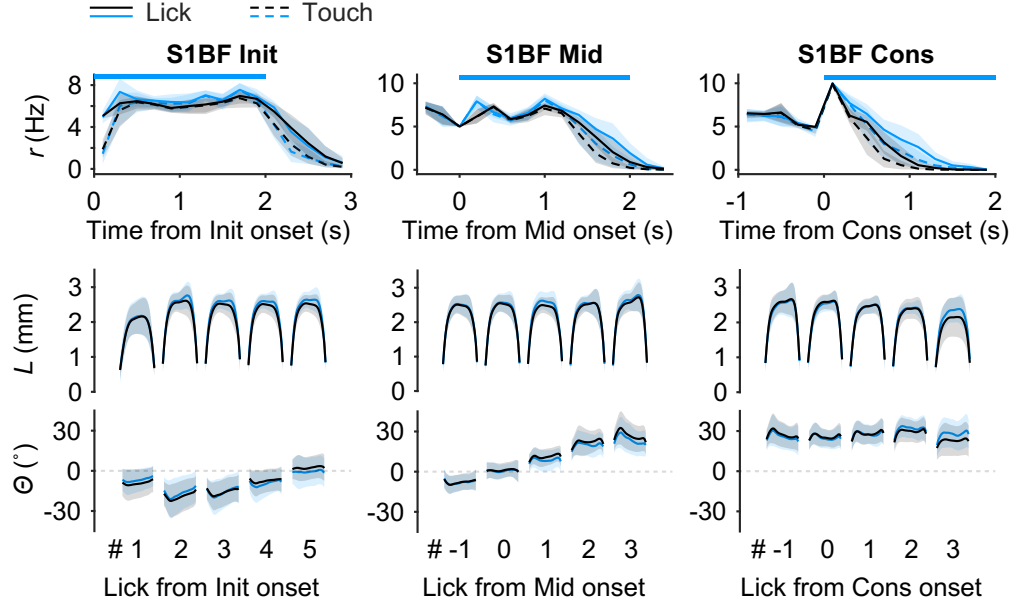
Figure 3.5: Effects of optoinhibition to S1BF

Effects of S1BF inhibition for each of the three periods. Top row, the rate of licks (solid lines) or touches (dashed lines) as a function of time for trials with (blue) or without (black) inhibition (mean ± 95% hierarchical bootstrap confidence interval). Horizontal bars indicate the 2 s long illumination. Time zero is the onset of inhibition, or the equivalent time in trials without inhibition. Middle and bottom rows, L and patterns (mean ± SD) for five consecutive licks in trials with (blue) or without (black) inhibition. Lick #0 is the lick that triggered inhibition, #1 the first lick after the trigger, #-1 the lick before, etc. Licks from inhibited trials are colored blue and those without inhibition are in black. n = 1399 trials for M1 from 5 sessions, one session per mouse.

3.3b). In 3 of the 5 mice, licks were largely absent (Figure 3.7a and 3.6a). Inhibiting S1TJ caused more moderate suppression (Figure 3.3a), and there was no obvious change when inhibiting other regions (Figure 3.4, 3.5, 3.6a, and 3.7a). When applied at mid-sequence, ALM inhibition also suppressed the production of licks, although less strongly. Inhibiting any other region at mid-sequence showed little effect.
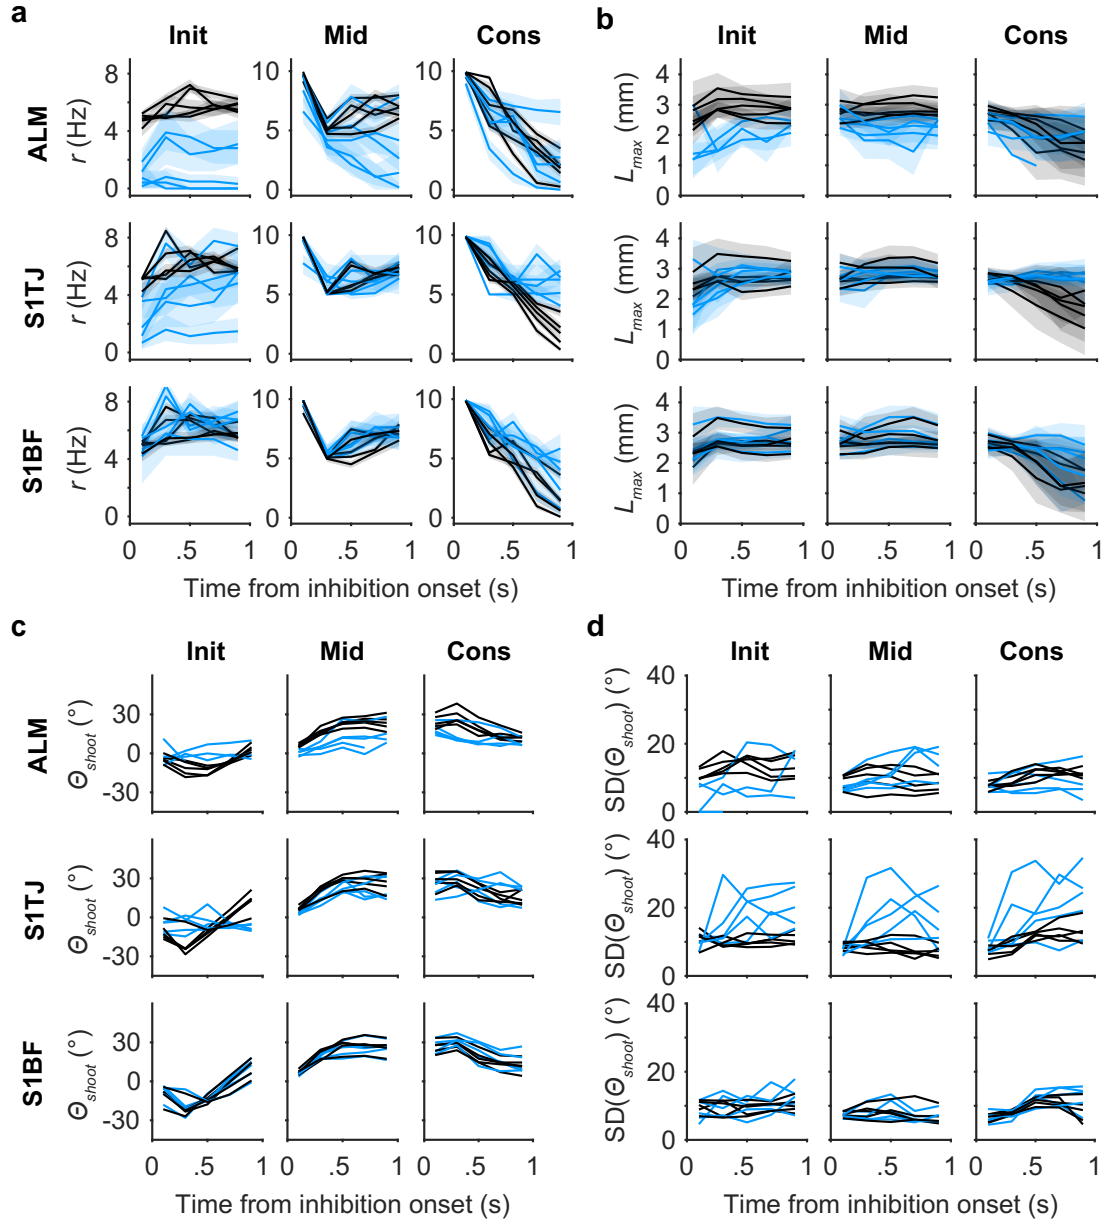
Figure 3.6: Effects of optoinhibition by individual mice

(a) Rate of licks (r) of individual mice (mean ± 95% bootstrap confidence interval) with (blue) or without (black) inhibition in ALM (top), S1TJ (middle) or S1BF (bottom) and in each of the three periods.

(b) Similar to (a) but quantifies L_{max} as a function of time for individual mice (mean ± SD).

(c) Similar to (b) but quantifies effects in Θ_{shoot}. The standard deviation of Θ_{shoot} ($SD(\Theta_{shoot})$) is separately plotted in (d) for clarity.

61
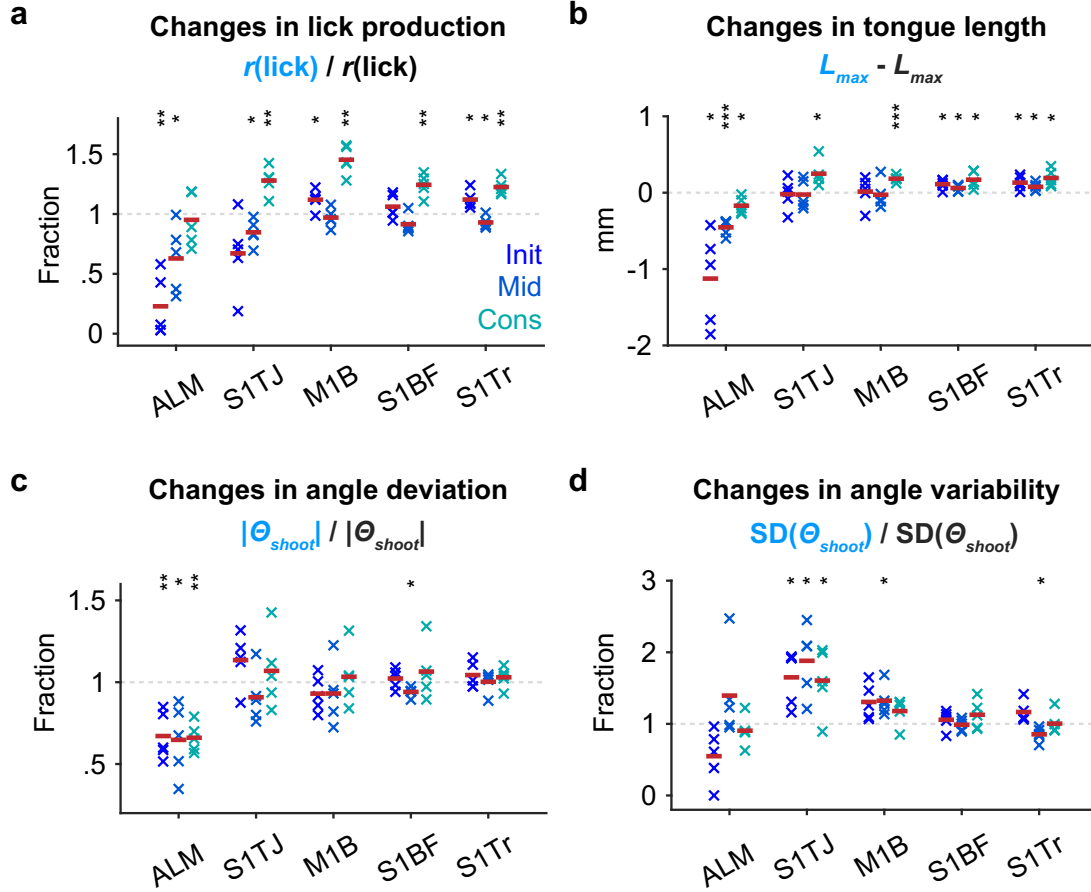
Figure 3.7: Summary of optoinhibition effects

(a) Changes in lick production across regions and periods, quantified as lick rate in trials with inhibition as a fraction of lick rate in trials without inhibition. Red bars indicate means (n = 5 mice). $*$ p < 0.05, $**$ p < 0.01, $***$ p < 0.001, not significant otherwise, two-tailed t-test.

(b) Same as **(a)** but for changes in L_{max}.

(c) Same as **(a)** but for changes in the ability to direct licks away from the midline, quantified using $|\Theta_{shoot}|$.

(d) Same as **(a)** but for changes in the variability of lick angles, quantified using the standard deviation of Θ_{shoot}.

3.3.3 Anterior Cortex Activity is Collectively Required for Sequence Termination

Previous studies have attributed M1B almost exclusively to body and limb control. To our surprise, when inhibiting M1B at water consumption, mice were impaired at stopping ongoing sequences (Figure 3.4a, top right). This prolonged licking was not due to additional attempts to reach the port for water as mice continuously made successful contacts. However, we also observed a trend toward an increase of lick production when inhibiting S1TJ, and more subtly, S1BF and S1Tr.

To test the possibility that inhibition of multiple areas caused persistent lick bouts due simply to spread of inhibition into M1B, we repeated all the above experiments with half of the illumination power (2 mW). The effects of ALM inhibition on sequence initiation, tongue length and angle control, and of S1TJ inhibition on angle control, remained consistent with, though much weaker than, our previous results using higher power (Figure 3.9). At consumption, however, inhibiting M1B resulted in similar deficit in terminating ongoing sequences. Surprisingly, this paradoxical effect also became evident when inhibiting ALM and S1TJ (Figure 3.8). Therefore, our observation that inhibition of multiple anterior cortical areas produces a deficit in sequence termination was not due to spread of inhibition into M1B.
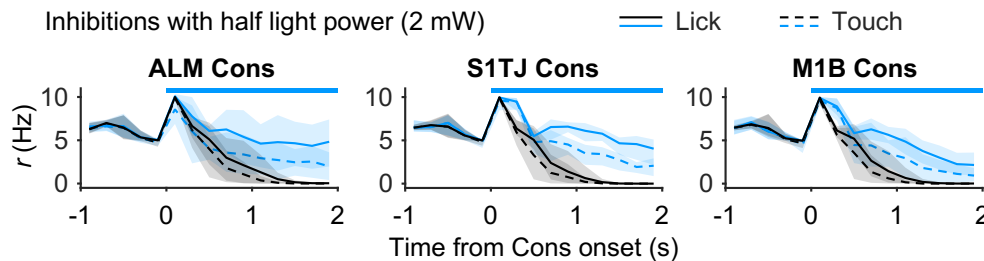
Figure 3.8: Effects of low power optoinhibition at consumption periods

Effects of low power (2 mW) inhibition at ALM (n = 850 trials), S1TJ (n = 774) and M1B (n = 784) during consumption period. The rate of licks (solid lines) or touches (dashed lines) as a function of time for trials with (blue) or without (black) inhibition (mean ± 95% hierarchical bootstrap confidence interval). Horizontal bars indicate the 2 s long illumination. Time zero is the onset of inhibition, or the equivalent time in trials without inhibition.
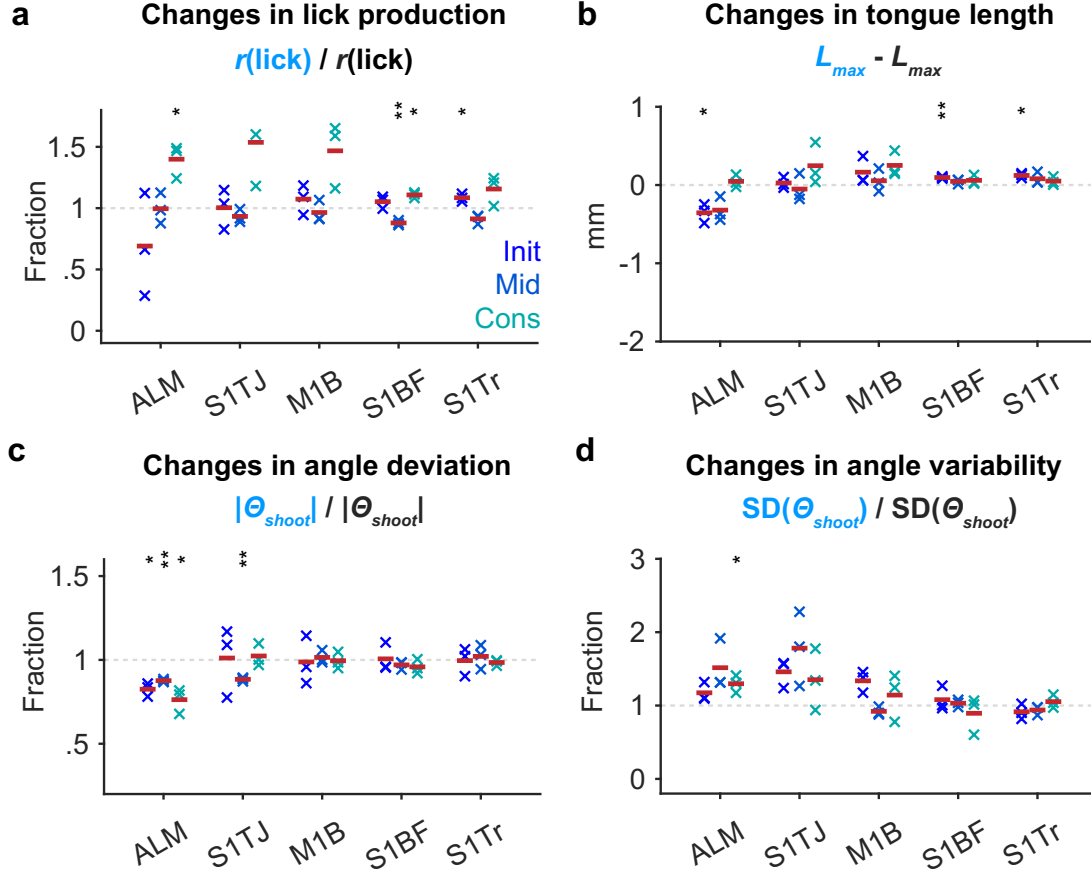
Figure 3.9: Summary of low power optoinhibition effects

(a) Changes in lick production across regions and periods, quantified as lick rate in trials with inhibition as a fraction of lick rate in trials without inhibition. Red bars indicate means (n = 3 mice). $*$ $p < 0.05$, $**$ $p < 0.01$, $***$ $p < 0.001$, not significant otherwise, two-tailed t-test.

(b) Same as (**a**) but for changes in L_{max}.

(c) Same as (**a**) but for changes in the ability to direct licks away from the midline, quantified using $|\Theta_{shoot}|$.

(d) Same as (**a**) but for changes in the variability of lick angles, quantified using the standard deviation of Θ_{shoot}.

Rather, our results indicate that sequence termination is an active process mediated collectively by multiple regions in the anterior cortex. Since our data show that ALM and S1TJ played active roles in tongue control, inhibiting them with high power presumably impaired both sequence execution and termination. Low-power inhibition appeared to spare the control of execution more than that of termination, revealing deficits in the latter.

3.3.4 M1TJ Inhibition Impairs Motor Control of Licking

Functional mapping and anatomical tracing from previous studies have shown that M1TJ is involved in the control of tongue/jaw motions (Mayrhofer et al., 2019; Komiyama et al., 2010) and that it communicates with S1TJ (Mayrhofer et al., 2019). To test the role of M1TJ in sequence control, we performed additional inhibition experiments in a separate group of mice. M1TJ inhibition showed similar effects as ALM inhibition, including suppressed lick production, shortened tongue length and reduced angle modulation (Figure 3.10).
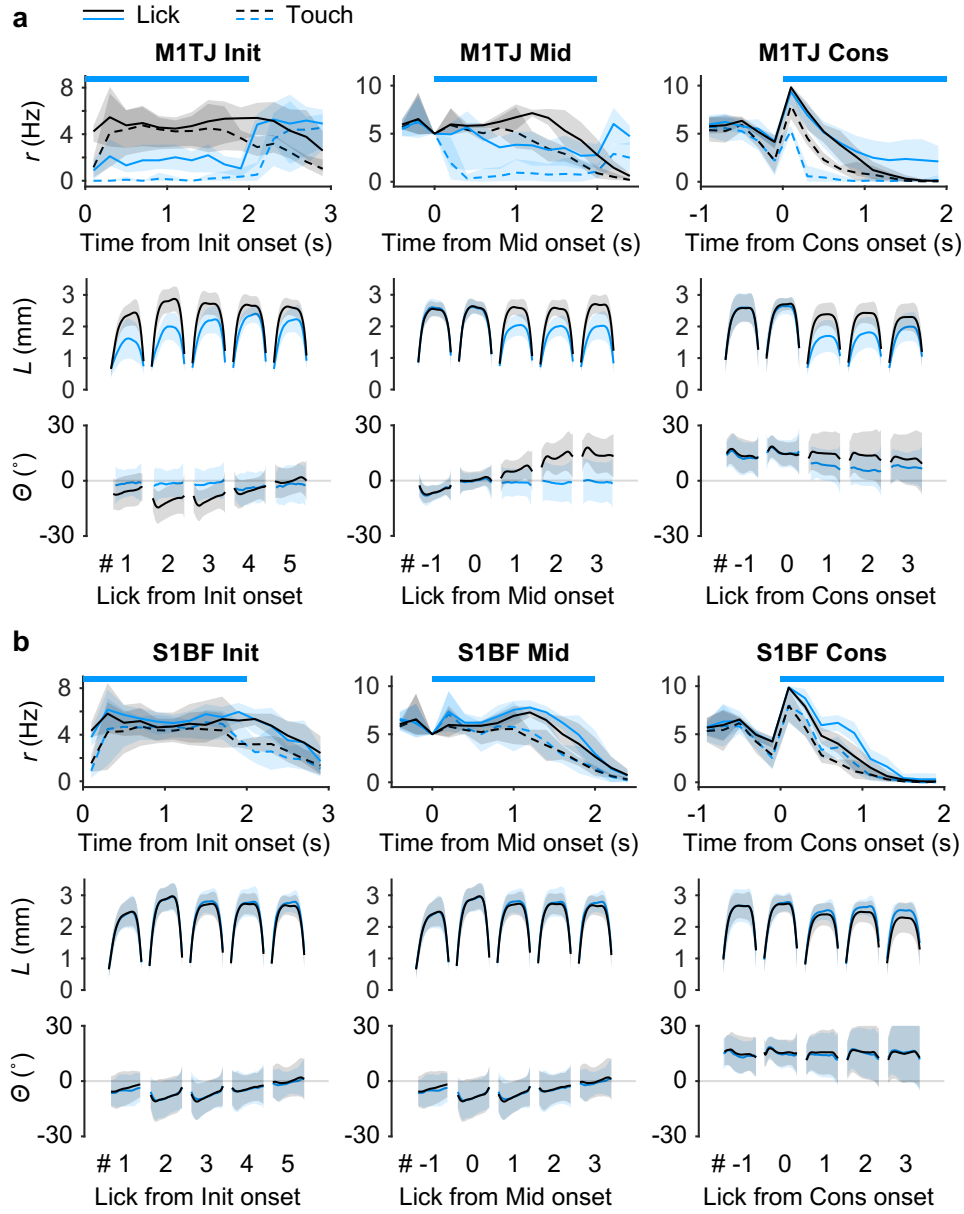
Figure 3.10: Effects of optoinhibition to M1TJ and S1BF

Effects of M1TJ (**a**) and S1BF (**b**) inhibition for each of the three periods. Top row, the rate of licks (solid lines) or touches (dashed lines) as a function of time for trials with (blue) or without (black) inhibition (mean ± 95% hierarchical bootstrap confidence interval). Horizontal bars indicate the 2 s long illumination. Time zero is the onset of inhibition, or the equivalent time in trials without inhibition. Middle and bottom rows, L and patterns (mean ± SD) for five consecutive licks in trials with (blue) or without (black) inhibition. Lick #0 is the lick that triggered inhibition, #1 the first lick after the trigger, #-1 the lick before, etc. Licks from inhibited trials are colored blue and those without inhibition are in black. n = 709 trials for M1TJ from 3 sessions, one session per mouse. n = 766 trials for S1BF.

3.4 Discussion

The role of rodent S1TJ during behavior has never been investigated before. To our surprise, S1TJ inhibition randomized the licking sequence while keeping the length of licks and the mice's ability to lick to the sides intact. There are at least two possible explanations for this effect. First, the inhibited sensory feedback is proprioceptive. From a feedback control point of view (Scott et al., 2015), a missing or perturbed proprioceptive feedback from the tongue may cause the motor output (such as lick angle in our case) to oscillate randomly (Fink et al., 2014). The lack of deficits in tongue length may be due to the fact that mice do not need to maintain tongue length at an intermediate level (like monkeys in Lin et al. (1993)). Therefore the length control largely operates in an opened-loop fashion. The second possibility is that inhibiting S1TJ abolishes the tactile feedback from the lick port. Without feeling the port, mice may start exploring around randomly. However, this second explanation seems less likely since mice sometimes miss the port spontaneously and they tend to keep licking at the expected port position instead of searching randomly. For future experiments, it would be interesting to see whether or not we can use S1TJ inhibition to mask out tactile feedback of the fourth lick in a standard sequence and trick mice to think backtracking has happened. If the mice lick back, it means the S1TJ inhibition indeed blocked tactile feedback from the port (though a negative result can have various interpretations).

Previous works have shown that inhibiting ALM can prevent reward-oriented voluntary licks (Komiyama et al., 2010). Our result generalizes its role from initiating simple rhythmic licking to initiating specific licking sequences. Furthermore, ALM is responsible for maintaining the proper length and angle modulation of the tongue throughout the course of the sequences. In contrast to a traditional view that consummatory licks are largely driven by brainstem reflexes to the sense of water (Miller, 2002) and independent from cortical control (Komiyama et al., 2010), we found that inhibiting ALM during consumption could dramatically impair mice's ability to direct licks to the side though only slightly affect the length of licks (Figure 3.3b). It is possible that the brainstem circuits are able to produce straight

67

consummatory licks in a "default" mode, but have no built-in mechanism to maintain a series of directed licks.

Our results show that sequence termination is an active process. If the consummatory licks at the end of each sequence are passively perpetuated by the sense of water, licking will automatically stop once water is fully consumed. However, when inhibiting any of the anterior cortical regions, we observed prolonged licking to the lick port even in the absence of water. The evidence for the neural control of sequence termination has largely focused on basal ganglia. Parkinson's and Huntington's patients have been reported to have deficits in stopping ongoing movement sequences (Benecke et al., 1987; Agostino et al., 1992). Many neurons in basal ganglia respond to sequence initiation and termination (Jin and Costa, 2010; Jin et al., 2014; Desrochers et al., 2015a). It is possible that the cerebral cortex is engaged in this process via the cortical-striatal loop, but may not be in a somatotopically specific way. More studies will be needed to dissect which neuronal subtypes (both in terms of genetic identity and connectivity) are mediating sequence termination and to track down the chain of signals.

The optoinhibition screening approach tries to strike a balance between experimental throughput and anatomical specificity. The spread of inhibition is not trivial to quantify (Li et al., 2019). Although the characterization shown in the Results only takes illumination power and corticocortical distance into account (Figure 3.2d), other factors such as the animal's behavioral state (anesthetized, awake, or engaged in task), the connectivity of different regions (e.g. inhibiting S1TJ may have less impact on ALM than inhibiting ALM on S1TJ), and even the sex- or age-dependent level of expression of opsins can all play a role in inhibition efficiency. A comprehensive titration is simply not practical. It is possible that the similarity in behavioral deficits caused by ALM or M1TJ inhibition may be in part due to the spread of light. However, the distance between ALM and S1TJ are far enough and we did not see activation of inhibitory neurons due to the spread of light. Furthermore, the behavioral deficits are qualitatively different in ALM and S1TJ. Future works will likely

transition from screening types of experiments to specific hypothesis-driven manipulations which can take advantage of, for example, virus-mediated local expression of opsins in a projection-defined subpopulation of neurons.

Chapter 4

Single-neuron and Population Coding of Sensorimotor Sequences across Cortical Area

4.1 Introduction

Single-neuron responses to the execution of motor sequences have been well-characterized in non-human primates. To briefly recapitulate what has been reviewed in Chapter 1.2.1, neurons in the primary motor cortex (M1) predominantly respond to specific movements regardless of the sequence context. In contrast, supplementary motor area (SMA), supplementary eye field (SEF), premotor cortex (PM), dorsolateral prefrontal cortex (dlPFC), and anterior cingulate cortex (ACC) all contain neurons, to varying degrees, that are tuned to a conjunction of specific movement, a movement's temporal position in a sequence (i.e. sequence rank), and/or the identity of neighboring movements. Although a functional cortical hierarchy has become a canonical view for visual processing (Felleman and Van Essen, 1991), it remains unclear what are the relationships among the various cortical regions for sensorimotor sequence control. Rodent cortex has fewer areal specializations (Preuss, 1995) and

may serve as a reduced model for us to better understand the essential processing principles. Unfortunately, we know very little about how neurons in rodent brains encode for sensori-motor sequences, let alone the underlying circuit basis. By establishing the novel sequence licking task (Chapter 2) and a systematic screening for important cortical regions (Chapter 3), we are well-positioned to start looking into the neural code in these key regions during sequence performance.

With advances in recording and data processing techniques (Steinmetz et al., 2018), it is now easy to obtain many dozens or even hundreds of single-unit activities simultaneously. However, it poses challenges to analyze and make sense of the large array of heterogeneous responses. This problem can be dealt with at single-neuron level or at population level. A common approach at single-neuron level is to classify neurons into artificially defined categories, such as movement-responsive or sequence rank-responsive neurons. This method is straightforward but subject to authors' intended framework of interpretation. One may end up throwing the baby out with the bathwater by missing the important but unexpected neural dynamics (Shenoy et al., 2013). Alternative to single-neuron level analyses, the activity of neuronal populations can be treated as a whole. For example, one can ask about correlations between neural responses and behaviors if the population activity can be used to predict behavioral variables via statistical models (Musall et al., 2019). Although being more quantitative, it shares a similar spirit to the artificial categorization of neuronal responses in the sense that both rely on prior hypotheses. On the other hand, unsupervised dimensionality reduction methods, such as principal component analysis (PCA) and its variants (e.g. Liu et al. (2014); Kobak et al. (2016); Williams et al. (2018)), are agnostic to experimental hypotheses and can be used to discover dominant dynamics in population activity. The present study combines both single-neuron and population level analyses to provide a consistent and complementary account of neural coding across cortical regions. Furthermore, we will take the best of both hypothesis-driven decoding and unsupervised analysis to reveal the nature of major dynamics in each region.

4.2 Methods

The use of mice, surgical procedures, behavioral task control, instrumentation, behavioral training procedures, and common analysis methods are described in Chapter 1.5.

4.2.1 Electrophysiology

Two types of silicon probe were used to record extracellular potentials. One (H3, Cambridge Neurotech) had a single shank with 64 electrodes evenly spaced at 20 µm intervals. The other (H2, Cambridge Neurotech) had two shanks separated by 250 µm, where each shank had 32 electrodes evenly spaced with 25 µm intervals. Before each insertion, the tips of the silicon probe were dipped in either DiI (saturated) or DiD (5-10 mg/mL) ethanol solution and allowed to dry. Probe insertions were either vertical or at 40° from the vertical line depending on the anatomy of the recorded region and surgical accessibility. Once fully inserted, the brain was covered with a layer of 1.5% agarose and ACSF, and was left to settle for ∼10 minutes prior to recording. Based on the depth of the probe tip, the angle of penetration, and the position of these sites, the location of units could be determined. Units recorded outside the target structure were excluded from analysis.

Extracellular voltages were amplified and digitized at 30 kHz via an RHD2164 amplifier board and acquired by an RHD2000 system (Intan Technologies). No filtering was performed at the data acquisition stage. Kilosort (Pachitariu et al., 2016) was used for initial spike clustering. We configured Kilosort to highpass filter the input voltage time series at 300 Hz. The automatic clustering results were manually curated in Phy for putative single-unit isolation. We noticed a previously reported issue of Phy double counting a small fraction of spikes (with exact same timestamps) after manually merging certain clusters, thus duplicated spike times in a cluster were post-hoc fixed to keep only one.

Cluster quality was quantified using two metrics (Figure 4.1a-c,e). The first was the percentage of inter-spike intervals (ISI) violating the refractory period (RPV). We set 2.5
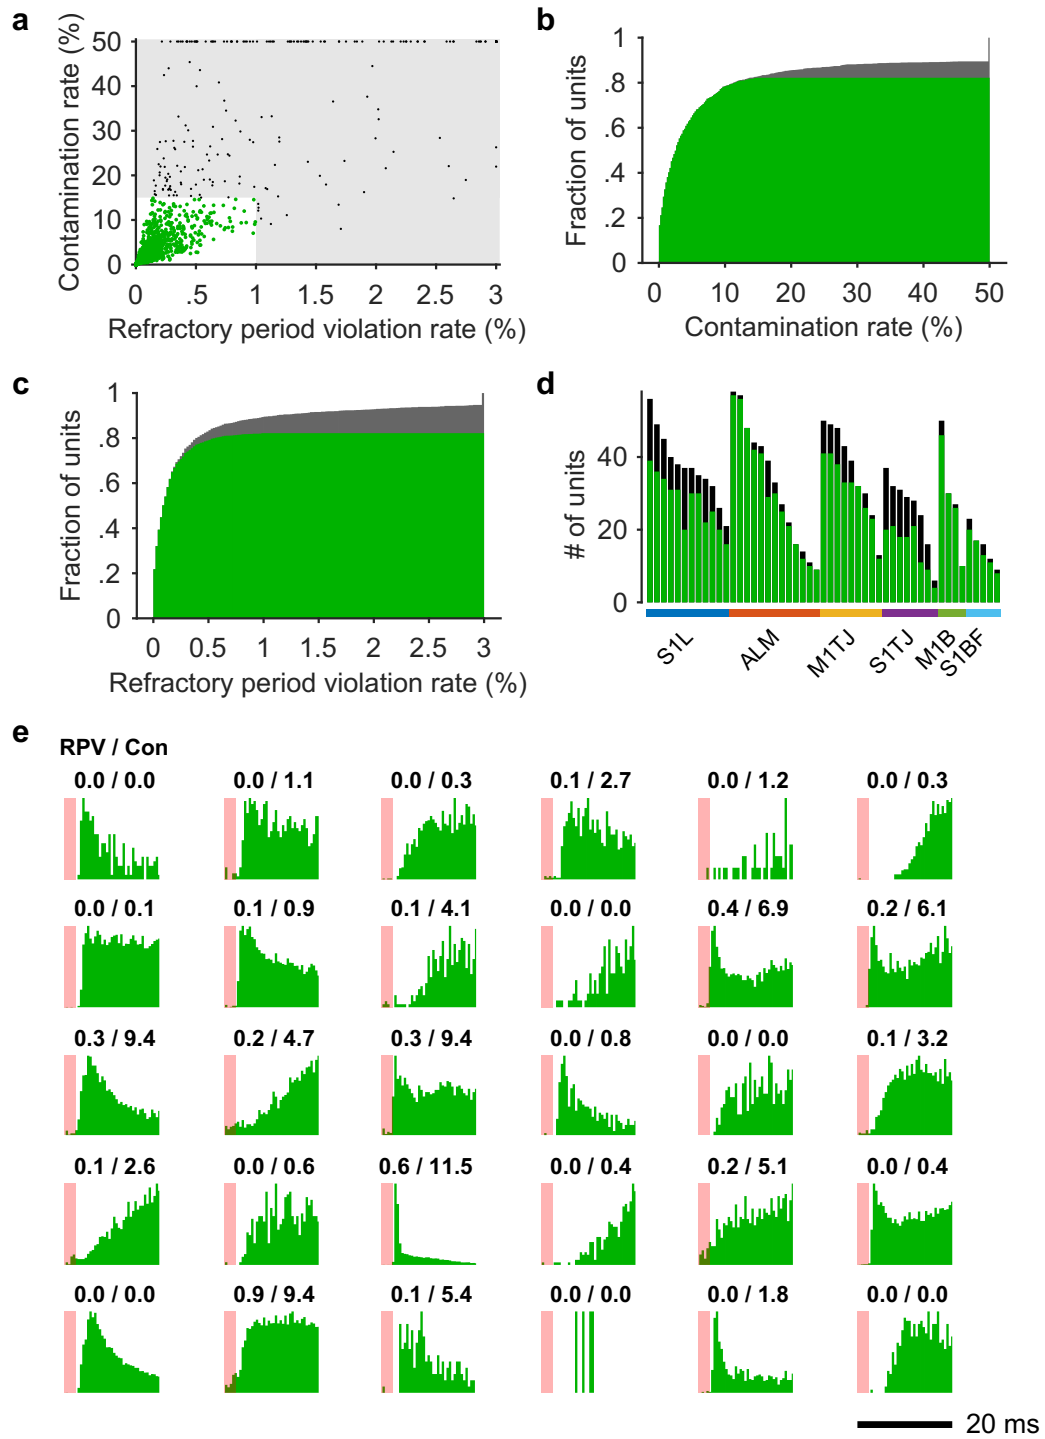
Figure 4.1: Quality control for spike sorting

Figure 4.1: Quality control for spike sorting

(a) Contamination rates and refractory period violation rates of all recorded single- (green) and multi-units (black). The shaded region shows the thresholds for assignment as multi- vs single-unit.

(b) CDF of contamination rate including single- (green) and multi-units (gray).

(c) Same as (**b**) but for refractory period violation rate.

(d) The number of single- (green) and multi-units (black) recorded in each session, grouped by brain area.

(e) ISI histograms of randomly selected single-units. Refractory period violation rates (RPV) and contamination rates (Con) are labeled on the top (in percent).

ms as the duration of the refractory period and used 1% as the RPV threshold above which clusters were regarded as multi-units. It has been argued that RPV does not represent an estimate of false alarm rate of contaminated spikes (Hill et al., 2011; Navratilova et al., 2016) since units with low spike rates tend to have lower RPV while high spike rate units tend to show higher RPV even if they are contaminated with the same percentage of false positive spikes. Therefore, we estimated the contamination rate based on the reported method (Hill et al., 2011). A modification was that we computed a cluster's mean spike rate from periods where the spike rate was greater than 0.5 spikes/s rather than from an entire recording session. As a result, the mean spike rate reflected more about neuronal excitability than task involvement. Any clusters with more than 15% contamination rate were regarded as multi-units. Combining these two criteria in fact classified less single-units than using a single, though more stringent, RPV of 0.5%. A low RPV can fail potentially well isolated fast spiking interneurons whose ISIs can frequently be shorter than the set threshold.

4.2.2 Histology

Mice were perfused transcardially with PBS followed by 4% PFA in 0.1 M PB. The tissue was fixed in 4% PFA at least overnight. The brain was then suspended in 4% agarose in PBS. A vibratome (HM 650V, Thermo Scientific) cut coronal sections of 100 m that were mounted and subsequently imaged on a fluorescence microscope (BX41, Olympus). Images

showing DiI and DiD fluorescence were collected in order to recover the location of silicon probe penetrations.

4.2.3   Trial selection and alignment

Due to individual variability, different mice tended to lick at slightly different rates within lick bouts. The same mice might also perform a bit faster in one sequence direction than the other. Even in a given direction, one might start faster and then slow down a little, or slower first and faster later. When aligning trials from heterogeneous sources, a 10% difference in lick rate, for instance, will result in a complete mismatch (reversed phase) of lick cycle after only 5 licks. Therefore, prior to the analyses that are sensitive to inconsistent lick rates, we linearly stretched or shrunk inter-lick intervals (ILIs) within each lick bout to a constant value of 0.154 s (i.e. 6.5 licks/s), which is around the overall mean. The timestamps of licks used to compute ILIs were contact onsets for touches or the time at L_{max} for missed licks. A lick bout was operationally defined as a series of consecutive licks in which every ILI must be shorter than $1.5 \times$ the median of all ILIs in the entire behavioral session. For ease of programming, we compensatorily scaled the time between the last lick of a trial and the start of the next trial to maintain an unchanged global trial time. Original time series, including spike rates and L', were obtained prior to temporal alignment. After alignment, the behavioral and neural time series were resampled uniformly at 400 samples/s.

After temporal alignment, we used a custom clustering algorithm to find a group of trials with the most similar sequence performance. First, all trials of a given sequence in a behavioral session were collected and a time window of interest was determined. For sigle-neuron analyses, we used the same time windows as the corresponding peri-event time histograms (PETH). For analyses related to population decoding during sequence execution, we used -1 to 1 s from the 4th (middle) touch. For analyses during consumption period, we used -0.5 to 1 s from the first lick touching water. The duration of a time window is denoted as T. Second, for each trial, a feature vector was constructed which included $7 \times T$

lick onset times and $7 \times T$ touch onset times within the time window. Insufficient timestamps were padded with zeros. Finally, pairwise euclidean distances were computed among feature vectors of all candidate trials and we chose a subset of N trials with the lowest average pairwise distance, i.e. those that have the most similar lick and touch times. The number N was set to 1/3 of available candidate trials with a minimal limit of $N = 10$ trials. We used this relatively low fraction mainly to handle the greater behavioral variability in sequences with backtracking. To handle trial-to-trial variability in sequence initiation time (defined as the interval from the cue onset to first touch onset), which was not captured in our feature vectors, prior to clustering we limited trials to those with sequence initiation time less than 1 second.

4.2.4   PETH, NNMF and t-SNE

Spike rates were computed by temporal binning (bin size: 2.5 ms) of spike times followed by smoothing (15 ms SD Gaussian kernel). The smooth PETHs were computed by averaging spike rates across trials. Each unit has 6 PETHs: 3 time windows (for sequence initiation, mid-sequence and sequence termination) each in 2 standard sequences (left to right and right to left). We excluded inactive units whose maximal spike rate across the 6 PETHs was less than 10 spikes/s. For the rest, we normalized PETHs of each unit to this maximal spike rate. To construct inputs to NNMF and t-SNE, the 6 PETHs of each unit were downsampled from 2.5 ms per sample to 25 ms per sample and were concatenated along time to form a single feature vector.

NNMF was performed using the MATLAB function "nnmf" with default options. We empirically chose to compute 10 clusters as too few clusters tend to merge response patterns tuned to adjacent stages of sequences, whereas too many clusters provided little help in extracting the major response patterns from the data. NNMF is a close relative of principal component analysis (PCA) that has gained increasing popularity for processing neural data. The algorithm finds a small number of activity patterns (equivalent to principal components

in PCA) along with a set of weights for each neuron, so that the original PETHs can be best reconstructed by weighted sums of those activity patterns. As a result, a small number of activity patterns (or dimensions) is usually able to capture the main structure of the original PETHs, and a neuron's weights quantifies the degree to which its activity reflects each pattern. In the context of clustering, each pattern describes representative activity of a cluster, and the pattern with the greatest weight for a neuron determines its cluster membership (Figure 4.9).

t-SNE is a nonlinear dimensionality reduction method. We used the MATLAB "tsne" function with customized options ('Algorithm', 'exact', 'Distance', 'cosine', 'NumDimensions', 2, 'Perplexity', 50, 'Standardize', false) to find a 2-D embedding of the feature vectors (Figure 4.8a). t-SNE per se does not cluster data. Therefore, we fitted Gaussian mixture models with various numbers of components and used BIC to determine the optimal number of clusters (Figure 4.8b). Since the t-SNE results can be sensitive to initial conditions, we repeated the computation with 50 different seeds of a random number generator and obtained a median cluster number of 8 (Figure 4.8c). The results shown are from the first run with 8 clusters.

4.2.5 Linear model and PCA

In each linear regression, the predictors were normalized spike rates of simultaneously recorded units and the response was one of the five behavioral variables (L, L', θ, D/TP or τ). Predictors and responses were sampled at 400 Hz and were not averaged across trials. PCA was performed using the same normalized spike rates.

Both single- and multi-units were included. To obtain normalized spike rates, we divided the original spike rate by the maximum spike rate or 5 Hz, whichever was greater. We adopted this "soft" normalization technique (Russo et al., 2018) to prevent weakly firing units from contributing as much variance as actively firing units. Note that this normalization was only necessary for PCA and did not affect the goodness of fit, R^2, of linear models.

L, L' and θ were directly available at 400 samples/s. However, these variables had values only when the tongue was outside of the mouth. Therefore, samples without observed values were either set to zero (for L) or excluded from regression (for L' and θ). D was defined as 1 if the sequence was from right to left and 2 if left to right. τ simply took sample timestamps as its values. In Chapter 5, TP was the same as D but defined based on the upcoming sequence.

Predicting single responses with dozens of predictors is prone to overfitting. Therefore, we chose the elastic-net variant of linear regression (using MATLAB function "lasso" with 'Alpha' set to 0.1), which penalizes big coefficients for redundant or uninformative predictors. A parameter λ controls the strength of this penalty. To find the best λ, we configured the "lasso" function to compute 10-fold cross-validated mean squared error (cvMSE) of the fit for a series of λ values. The smallest cvMSE indicates the best generalization, i.e. the least overfit. We conservatively chose the largest λ value such that the cvMSE was within one standard error of the minimum cvMSE. For each model, we derived the R^2 from this cvMSE.

Each of the 5 regressions resulted in a vector of coefficients comprising one coefficient for each unit, and a constant. Each coefficient vector and constant was used to predict or decode the corresponding behavioral variable from a vector (or, for multiple time points, a matrix) of population spike rates. We did not perform additional cross-validation in decoding because (1) 30% of the decoding for standard sequences (0.5 to 0.8 s in Chapter 4 and -1.3 to -1 s in Chapter 5) was from new data; (2) all decoding in backtracking sequences and during consumption period was from new data; and (3) the model has been proven the best generalization via cross-validation when selecting λ. The term "coefficient vector" and "coding axis" are used interchangeably.

The percent variance explained (VE) by principal components was simply derived from the singular values. To compute VE by each coding axis, we first obtained its unit vector and projected population spike rates onto it. The variance of the projected values is Var(explained). The total variance, Var(total), of the population activity is the sum of

78

variance of all units. Finally, VE equals Var(explained) / Var(total) × 100%.

4.2.6 Canonical correlation

In each session, we computed the trial-averaged neural trajectories in the five dimensional coding subspace and trajectories in the PC subspace from standard sequences (-0.5 to 0.8 s from 4th touch onset; right to left and left to right trajectories were concatenated). Canonical correlations were computed using MATLAB "canoncorr" function between trajectory matrices with the same number of dimensions. N correlation coefficients (r) quantified the correlation between the activity in each pair of the N dimensions after transformation. The average r across the N values reflected the overall alignment between the two transformed trajectories.

4.3 Results

4.3.1 Single-unit Responses Tile Sequence Progression

We used silicon probes to record from multiple brain regions from both hemispheres (Figure 4.2) during the task and obtained a total of 1312 single-units and 284 multi-units (Methods; Figure 4.1) from 51 recording sessions. Perievent time histograms (PETHs) of single-unit spiking (Figures 4.3 - 4.6) exhibited a wide variety of patterns prior, during, and after sequence execution.

To discover the main themes of single neuron level encoding in an unsupervised way, we pooled neurons from all brain regions and clustered them based on their PETHs using non-negative matrix factorization (NNMF, a technique closely related to principal component analysis, or PCA, and K-means clustering; Methods). We chose NNMF for its simplicity as a linear technique, its lack of an orthogonality constraint (present in PCA), and its duality as both a dimensionality reduction and clustering method.

We computed 10 clusters of PETHs (Figure 4.3 - 4.7). Cluster #1 neurons showed high
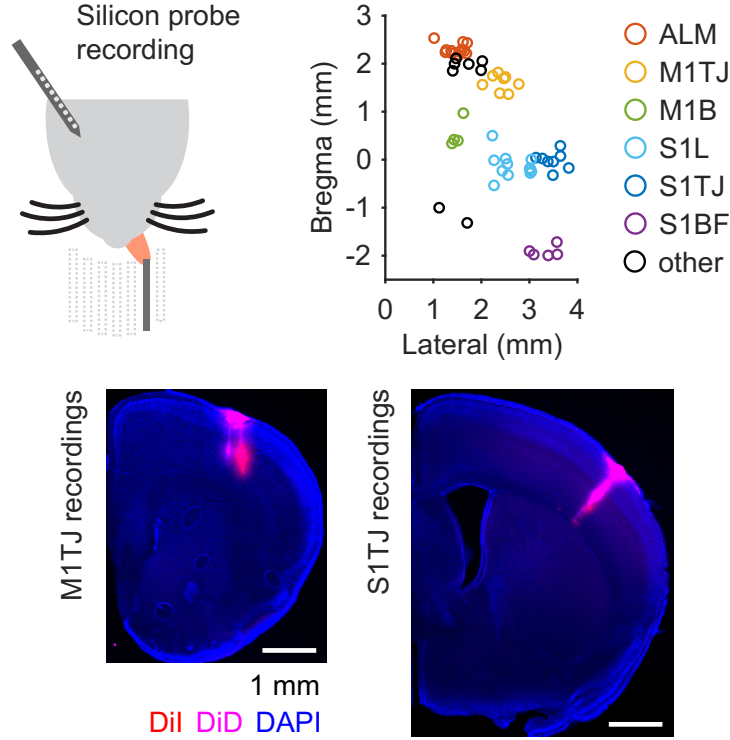
Figure 4.2: Silicon probe recording

Top left, silicon probe recording during the sequence licking task. Top right, histologically verified locations of silicon probe recordings. Recordings were made in both hemispheres but are illustrated for one. The plotted coordinates were randomly jittered by ± 0.05 mm to avoid visual overlap. Bottom, example sections showing the dye-labeled probe tracks. DiI and DiD were used in different penetrations.

activity outside sequence execution. Cluster #2 showed transient activation at sequence initiation. In contrast to the first two, all other clusters exhibited directional selectivity at specific stages, temporally tiling the behavioral sequences. Interestingly, neurons in clusters #9 and #10 reached peak activation when licking sequences stopped. Our finding of recurring PETH motifs was not due to our use of the specific NNMF method or the chosen number of clusters, as a different method yielded consistent results (Methods; Figure 4.8). Across NNMF clusters, activity of many neurons was smoothly modulated over hundreds of milliseconds; others, especially common in M1TJ and S1TJ, showed rapid modulation with individual licks. None of the clusters was biased to represent a small minority of neurons (Figure 4.9). However, the different cortical regions harbored different proportions of the 10
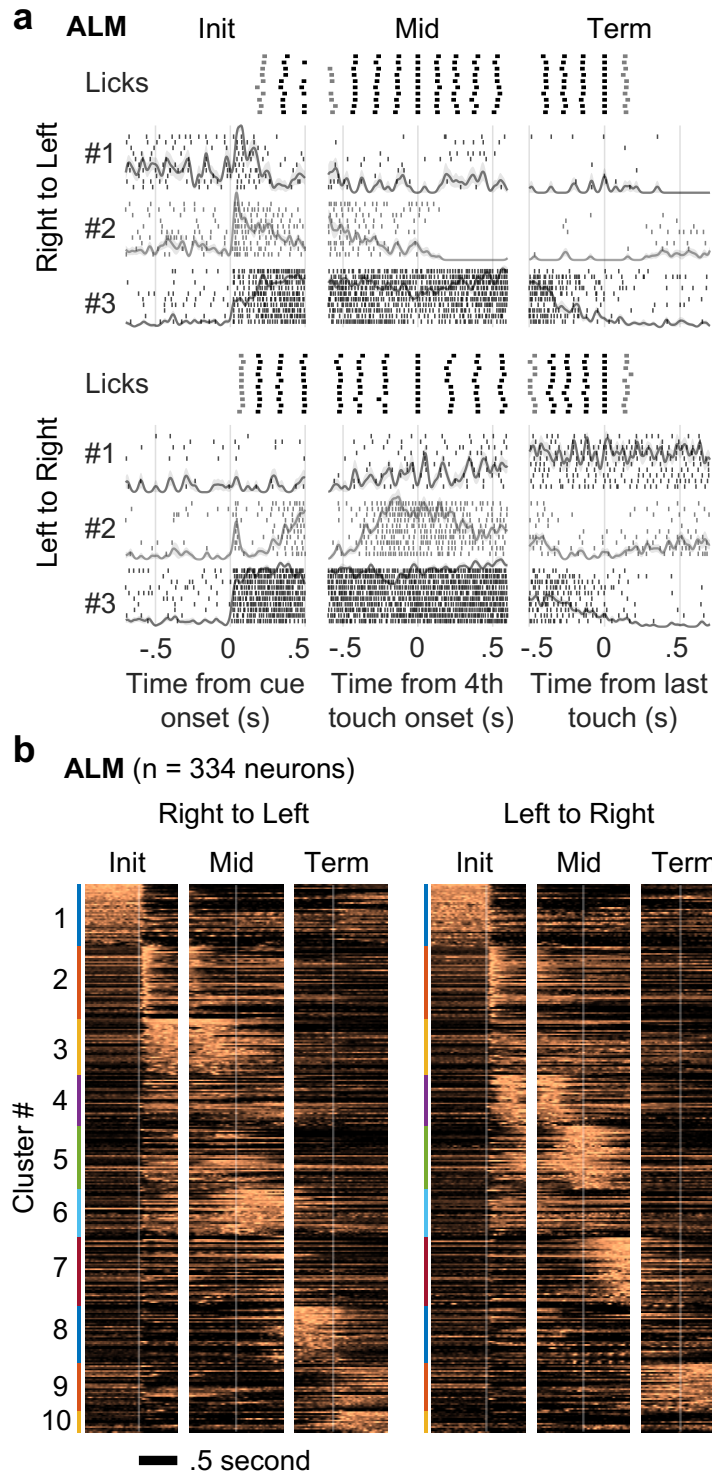
Figure 4.3: Responses of ALM neurons during behavioral sequences
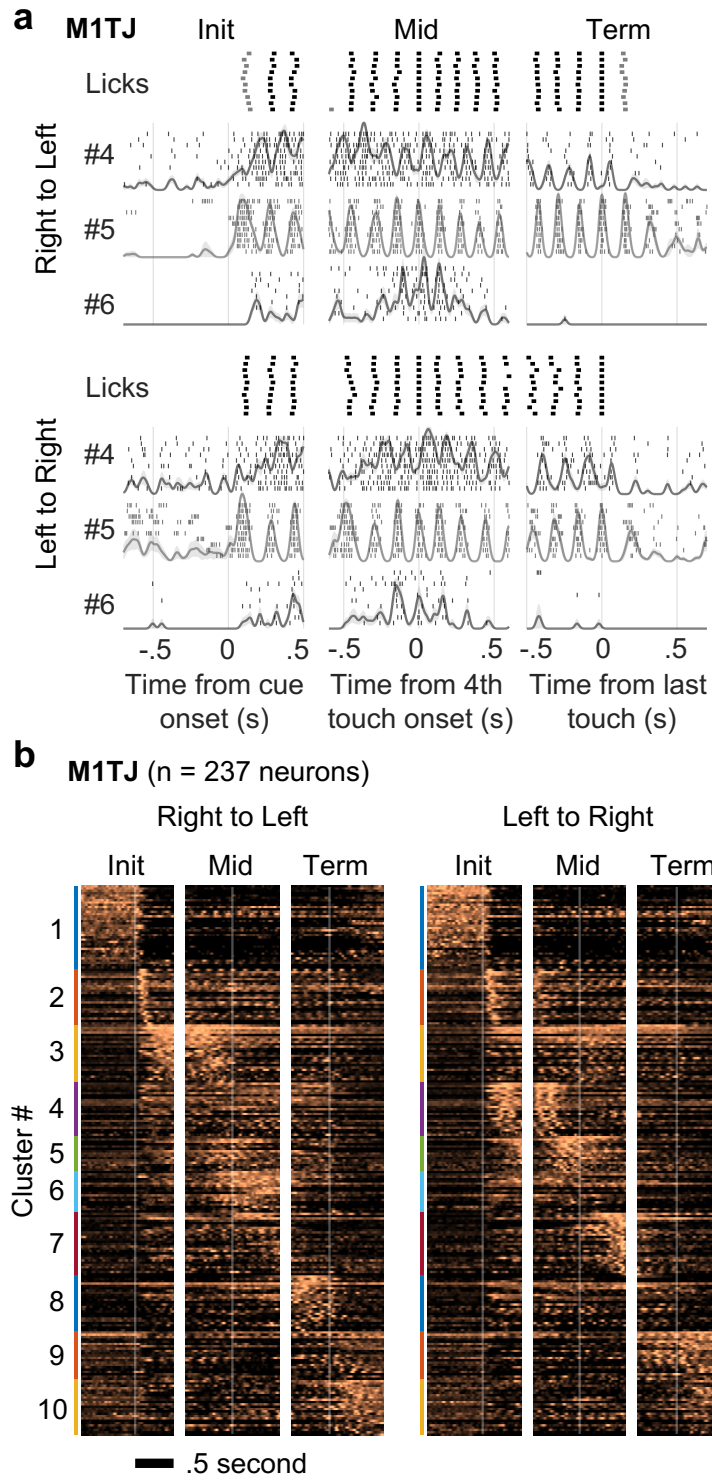
81

Figure 4.4: Responses of M1TJ neurons during behavioral sequences

a

S1TJ   Init   Mid   Term

Licks

Right to Left   #7   #8   #9

Licks

Left to Right   #7   #8   #9

-.5   0   .5   -.5   0   .5   -.5   0   .5

Time from cue onset (s)   Time from 4th touch onset (s)   Time from last touch (s)

b

S1TJ (n = 118 neurons)

Right to Left   Left to Right

Init   Mid   Term   Init   Mid   Term

Cluster #   1   2   3   4   5   6   7   8   9   10

.5 second
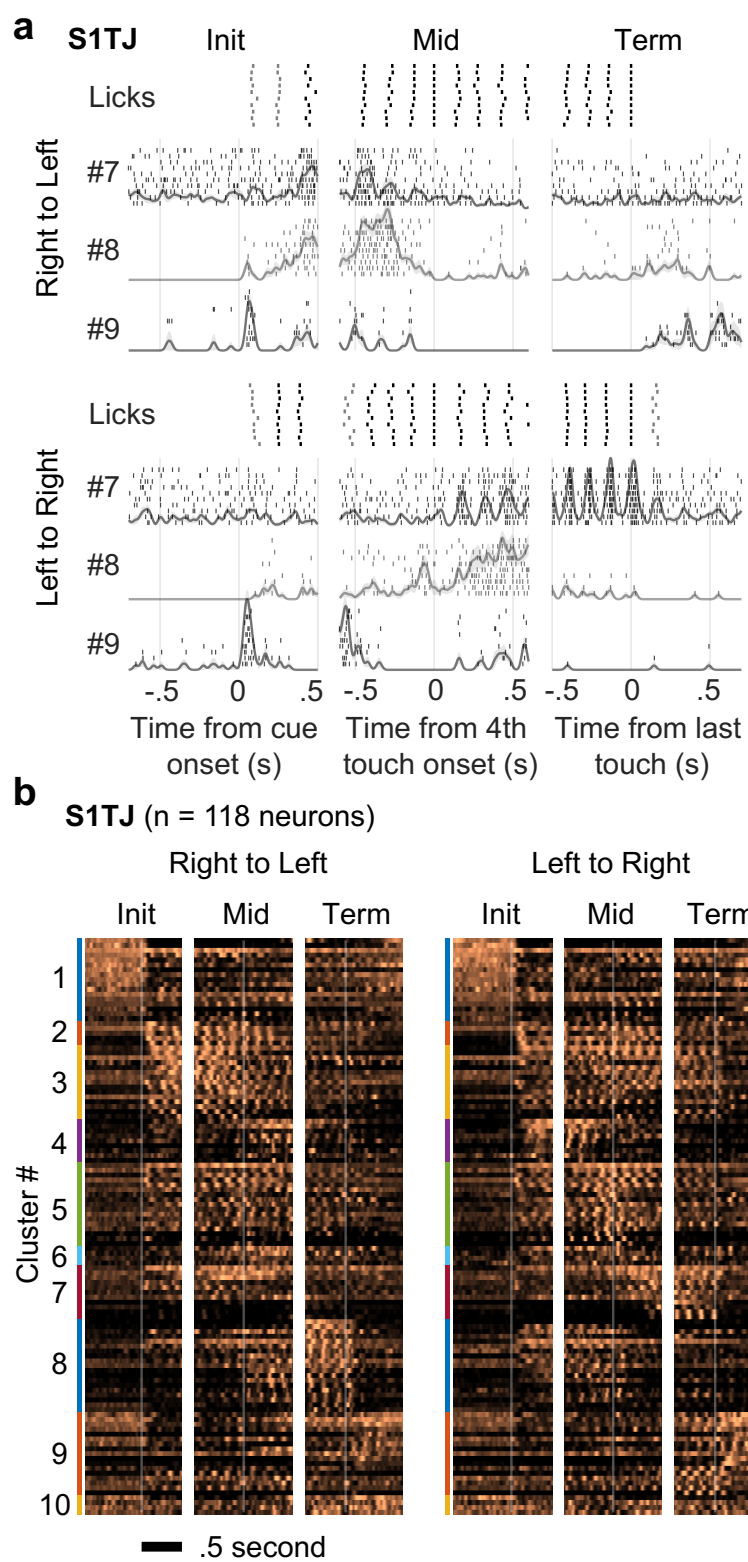
Figure 4.5: Responses of S1TJ neurons during behavioral sequences

Figure 4.3: Responses of ALM neurons during behavioral sequences

(a) Responses of three simultaneously recorded ALM neurons aligned with right to left (top) or left to right (bottom) licking sequences at initiation (left), mid-sequence (middle) and termination (right; defined as the onset of the last consummatory lick). For each sequence direction, the first row shows rasters of lick times (touches in black and misses in gray) from 10 selected trials (Methods). Stacked below are spike rasters and the corresponding PETHs from the same 10 trials for each example neuron.

(b) Normalized PETHs of all ALM neurons plotted as heatmaps. Neurons are grouped together by functional clusters (Results) and labeled by color bands. Time zero for each period is marked by the vertical lines and is consistent with (**a**).

Figure 4.4: Responses of M1TJ neurons during behavioral sequences

(a) Responses of three simultaneously recorded M1TJ neurons aligned with right to left (top) or left to right (bottom) licking sequences at initiation (left), mid-sequence (middle) and termination (right; defined as the onset of the last consummatory lick). For each sequence direction, the first row shows rasters of lick times (touches in black and misses in gray) from 10 selected trials (Methods). Stacked below are spike rasters and the corresponding PETHs from the same 10 trials for each example neuron.

(b) Normalized PETHs of all M1TJ neurons plotted as heatmaps. Neurons are grouped together by functional clusters (Results) and labeled by color bands. Time zero for each period is marked by the vertical lines and is consistent with (**a**).

Figure 4.5: Responses of S1TJ neurons during behavioral sequences

(a) Responses of three simultaneously recorded S1TJ neurons aligned with right to left (top) or left to right (bottom) licking sequences at initiation (left), mid-sequence (middle) and termination (right; defined as the onset of the last consummatory lick). For each sequence direction, the first row shows rasters of lick times (touches in black and misses in gray) from 10 selected trials (Methods). Stacked below are spike rasters and the corresponding PETHs from the same 10 trials for each example neuron.

(b) Normalized PETHs of all S1TJ neurons plotted as heatmaps. Neurons are grouped together by functional clusters (Results) and labeled by color bands. Time zero for each period is marked by the vertical lines and is consistent with (**a**).

S1BF (n = 55 neurons)
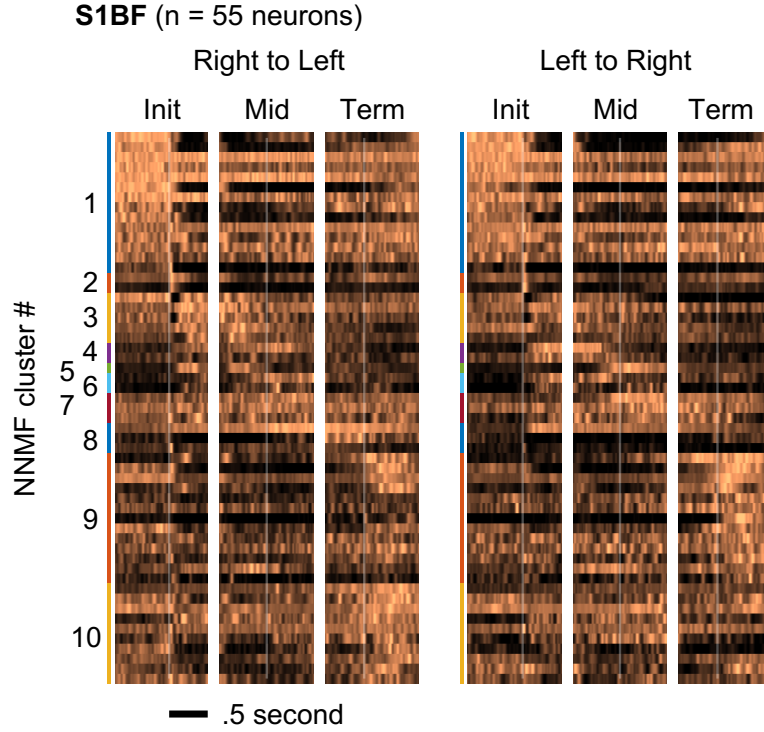
Figure 4.6: Responses of S1BF neurons during behavioral sequences

Normalized PETHs of all S1BF neurons plotted as heatmaps. Neurons are grouped together by functional clusters and labeled by color bands. Time zero for each period is marked by the vertical lines.

clusters. In contrast to ALM, M1TJ and S1TJ, the 7 clusters from #2 through #8 made up only 1/3 of all S1BF neurons (Figure 4.6). With neurons pooled from all regions, different cortical depths contained similar proportions of clusters (Figure 4.10a). Although the cluster proportions in individual regions could differ by cortical depth (e.g. cluster #2 units appeared less in S1TJ than in ALM or M1TJ), it is unclear whether or not this was due to an insufficient number of neurons once subdivided by depths and clusters (Figure 4.10b).

In sum, we observed both variety and commonality at the level of single-neuron responses. However, do patterns of activity arising from these single-unit responses encode behavioral variables important for sequence control?
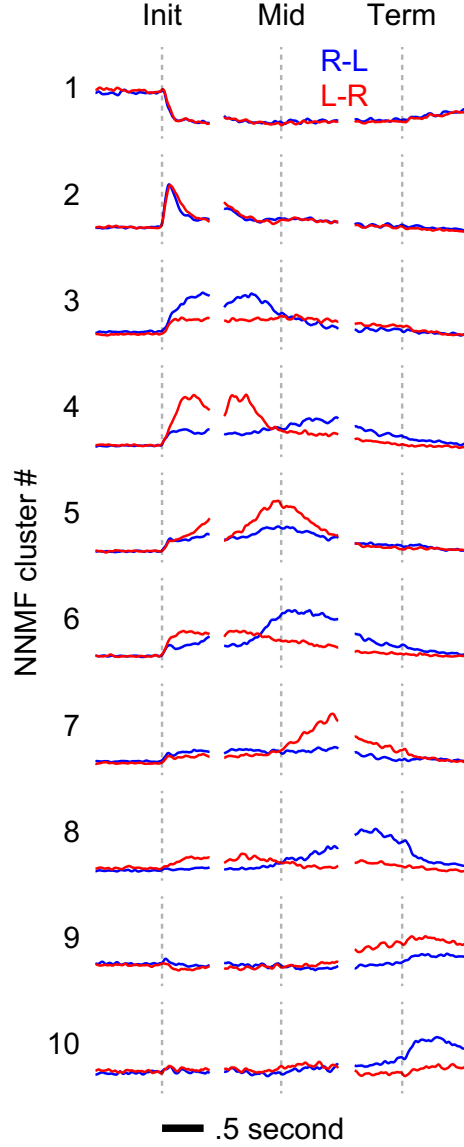
Figure 4.7: Mean responses of NNMF clusters

NNMF components that represent each of the ten PETH clusters. Right to left (blue) and left to right (red) activities (mean $\pm$ 95% bootstrap confidence interval) are overlaid together. The vertical lines are located at time zero in each period. The height of the lines represents the scale of normalized neuronal activity from 0 to 1.
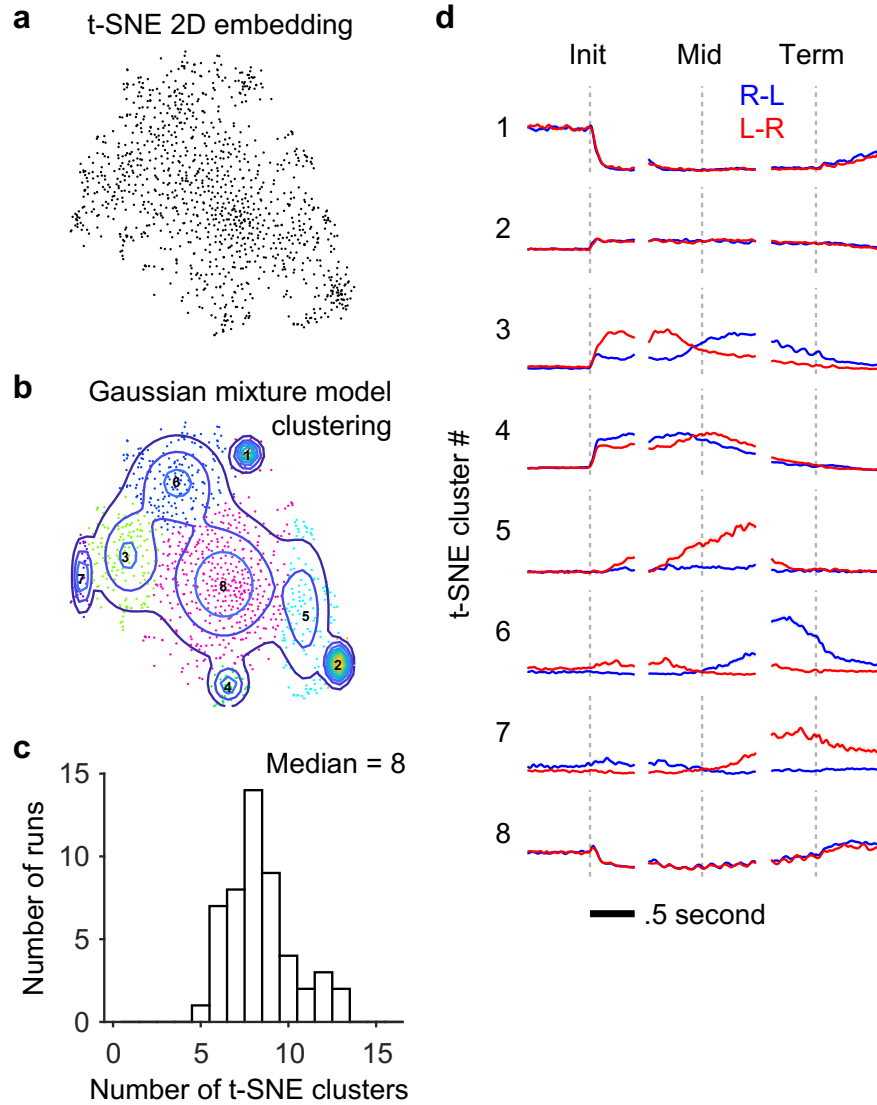
Figure 4.8: Clustering neuronal responses using t-SNE and Gaussian mixture models

(a) Two-dimensional embedding of neuronal PETHs via t-SNE.

(b) Nine clusters were fitted using a Gaussian mixture model.

(c) Distribution of the number of clusters computed with 50 different random number seeds.

(d) Mean responses for each of the ten PETH clusters. Right to left (blue) and left to right (red) activities (mean $\pm$ 95% bootstrap confidence interval) are overlaid together. The vertical lines are located at time zero in each period. The height of the lines represents the scale of normalized neuronal activity from 0 to 1.
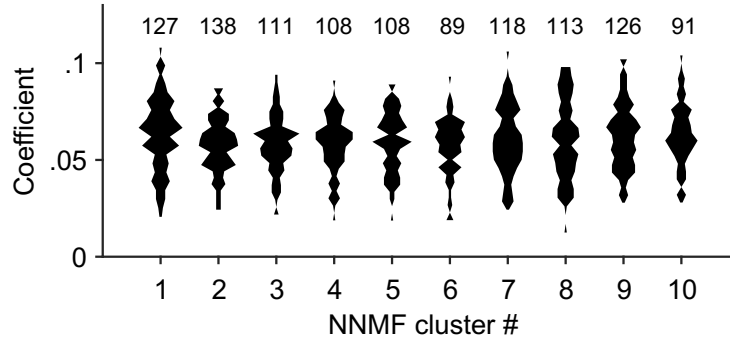
Figure 4.9: NNMF cluster membership

Distribution of membership coefficients for each of the NNMF clusters. The number of neurons in each cluster is labeled at the top.
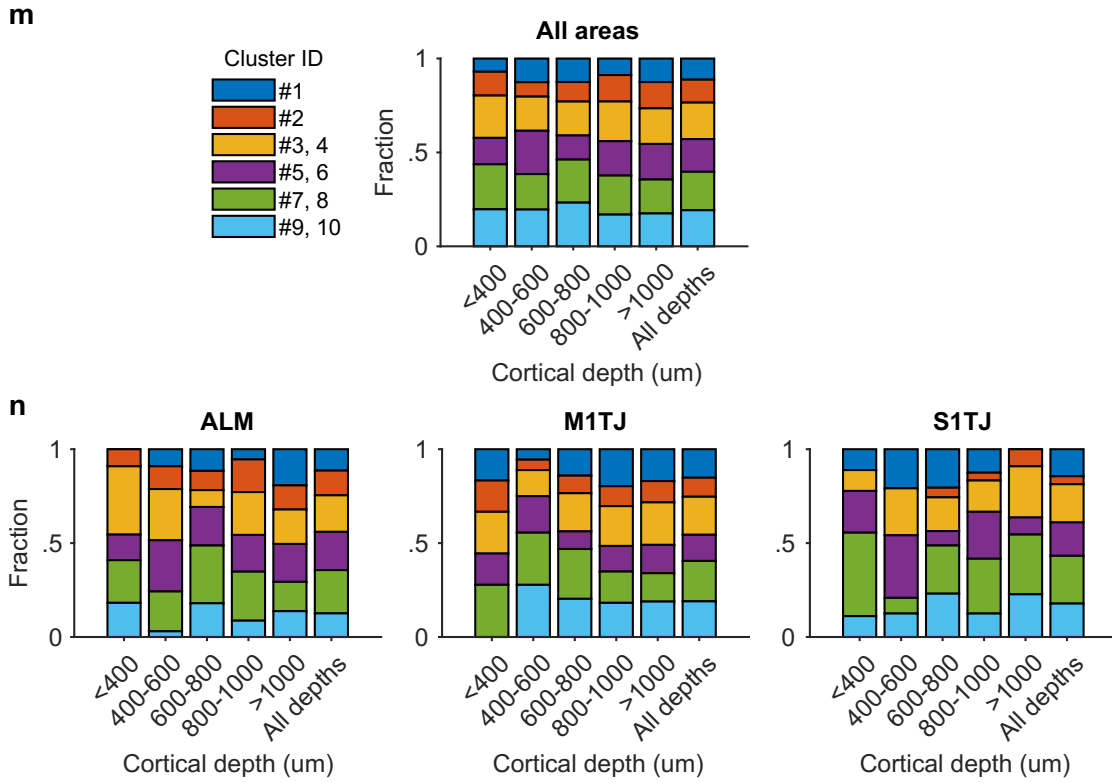
Figure 4.10: Distributions of functional clusters by cortical depth

(a) Proportions of neurons (n = 1312 in total) from different clusters at different cortical depths. Some clusters were grouped together (e.g. 3 and 4) since they fired around the same time but only differed in direction selectivity.

(b) Similar to (**a**) but broken down for ALM (n = 334), M1TJ (n = 237) and S1TJ (n = 118).

4.3.2 Hierarchical Population Coding of Behavioral Variables across Cortical Areas

In our sequence licking task, the brain needs to encode instantaneous tongue length (L) and angle (θ), presumably both for motor output and sensory feedback. The encoding of velocity (L') could also be used to indirectly control tongue position. However, encoding only instantaneous kinematics is not enough, as mice need to string individual licks into sequences. Sequences in one direction also require a different organization of licks compared with sequences in the other direction. One way to organize a sequence is to have a slowly varying encoding of expected target position, regardless of the presence or absence of a lick at a given moment in time. Another, mutually non-exclusive, way is to separately encode sequence direction (D) and relative sequence time (τ). The variable τ can also serve as a proxy for sequence progress or "distance to goal". The five behavioral variables, L, L', θ, D and τ, were all recorded (or derived) with a temporal resolution of 2.5 ms (Figure 4.14 and 4.15). Conveniently, any pair of these variables is uncorrelated (Figure 4.11). Therefore, being able to encode one is of little or no help to encode any other.

Linear models, though simple, are powerful methods to uncover information encoded by
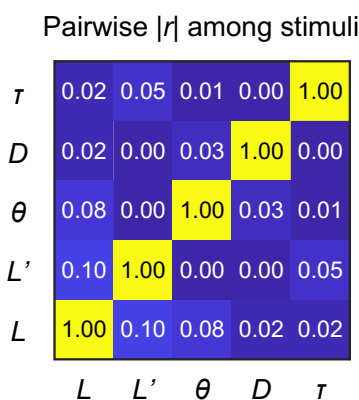
Figure 4.11: Behavioral variables are uncorrelated

Absolute mean pairwise Pearson's correlation coefficients among the five behavioral variables. (n = 35 sessions)
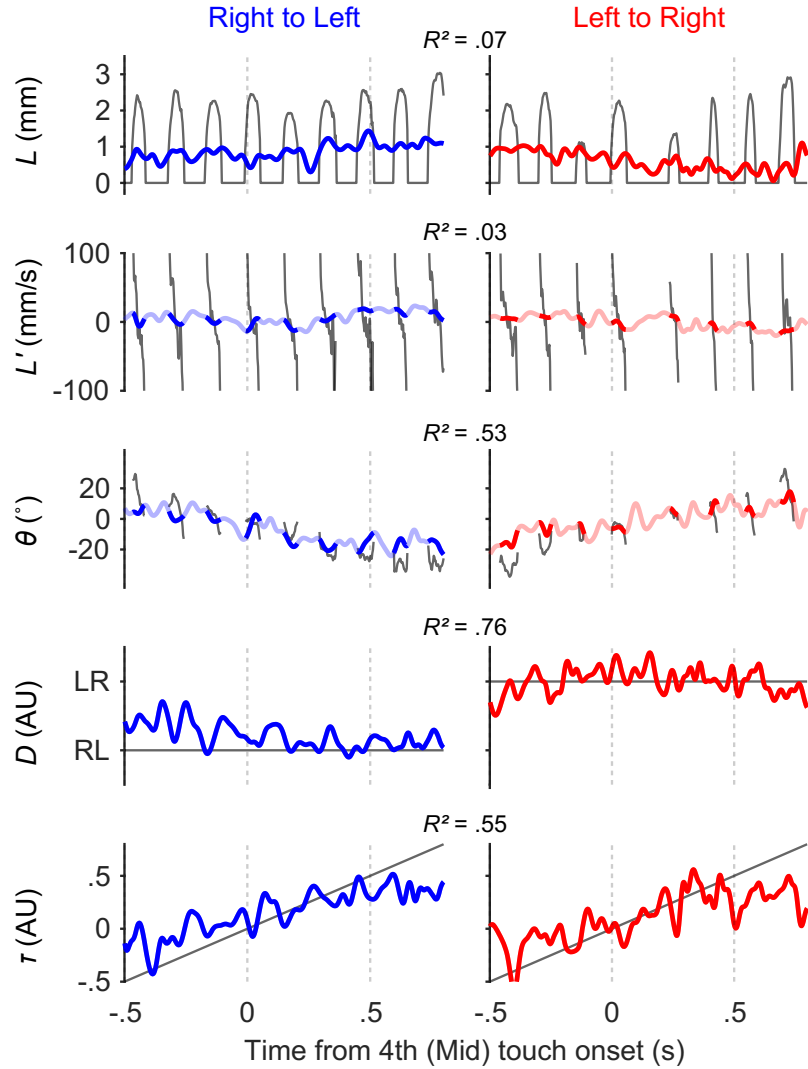
Figure 4.12: Single-trial decoding of behavioral variables in ALM

Single-trial decoding of the five behavioral variables (rows; black traces) from 58 simultaneously recorded ALM units in a right to left (left) and a left to right (right) sequence.

a population of neurons (Kriegeskorte and Douglas, 2019). For each recording session, we performed separate linear regressions (Methods) to obtain unit weights (and a constant) for each of the five behavioral variables, such that a weighted sum of spike rates from simultaneously recorded units (31 ± 14 units; mean $\pm$ SD) plus the constant best predicted the value of a behavioral variable. We used cross-validated R^2 values to quantify how well the recorded population of neurons encoded each behavioral variable. Importantly, we used only
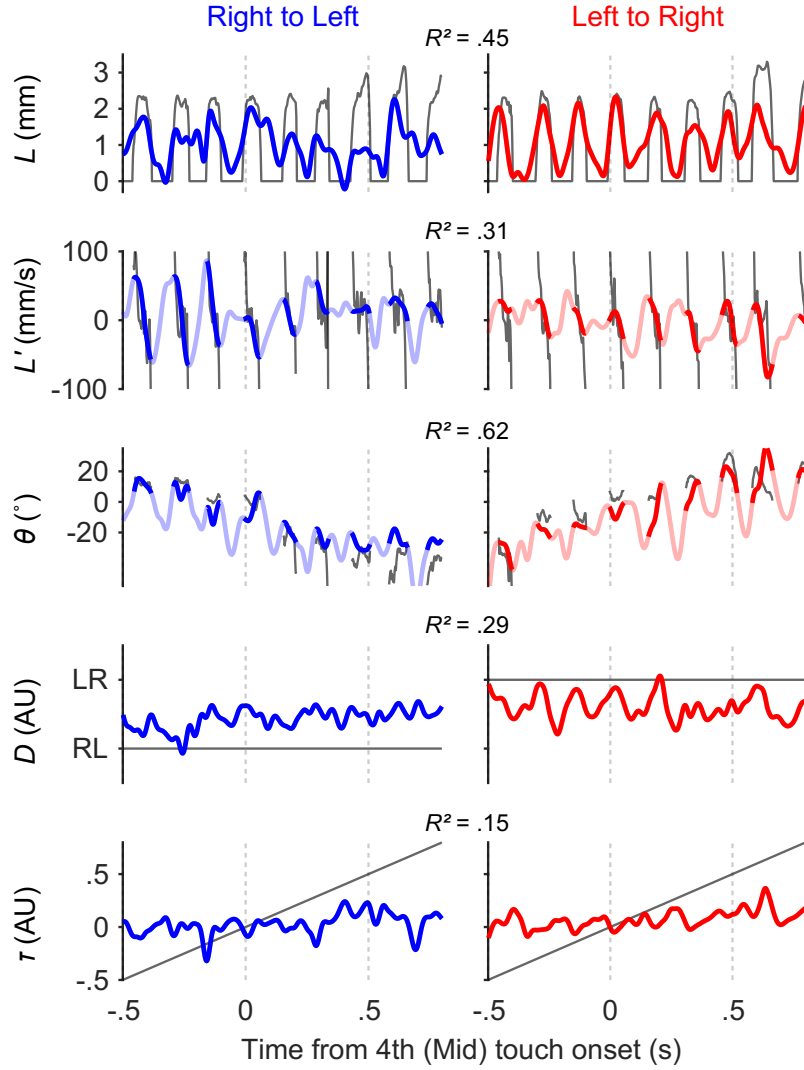
Figure 4.13: Single-trial decoding of behavioral variables in S1TJ

Single-trial decoding of the five behavioral variables (rows; black traces) from 29 simultaneously recorded S1TJ units in a right to left (left) and a left to right (right) sequence.

data from standard sequences to fit the models, but used these same models to decode both standard and backtracking sequences. The five behavioral variables could be decoded from population activity even for single trials (example from ALM is shown in Figure 4.12 and S1TJ in Figure 4.13).

Despite our observation that ALM inhibition uniquely produced a deficit in maintaining *L* during sequences, neuronal populations in this area showed weak encoding of instantaneous
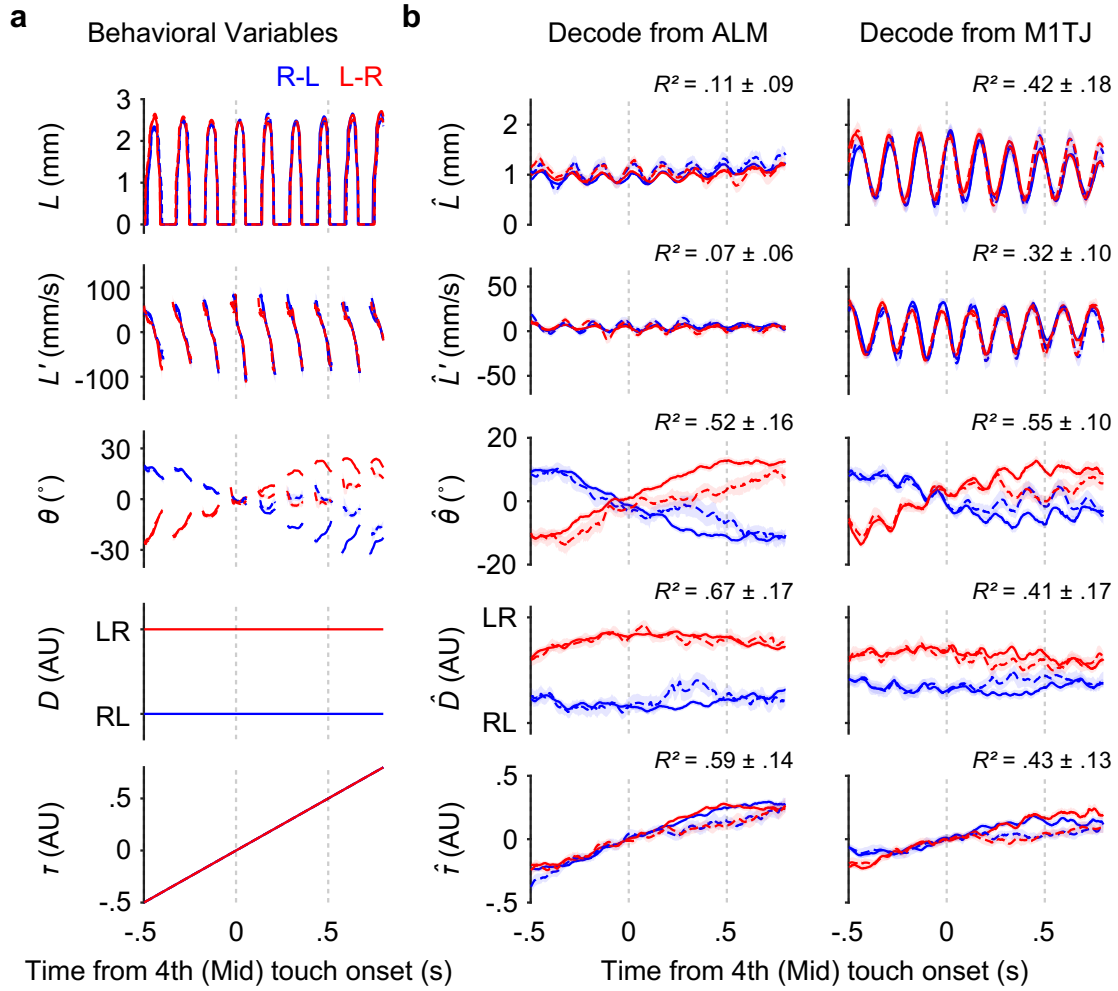
Figure 4.14: ALM and M1TJ population decoding of behavioral variables

(a) The five recorded (or derived) behavioral variables averaged across trials (n = 2684 trials; mean ± 99% bootstrap confidence interval) with standard sequences (solid lines) or backtracking sequences (dashed lines), each of which is either right to left (blue) or left to right (red). Time along the x-axis is aligned to the fourth lick in the sequence. Time points where more than 80% of trials did not have observations are not plotted.

(b) Decoding of the five behavioral variables (rows) from populations recorded in ALM and M1TJ (columns). Cross-validated R^2 for each region and variable is given (mean ± SD; ALM, n = 13 recordings; M1TJ, n = 9 recordings). Same plotting conventions as in **(a)**.

L and L' (Figure 4.14 and 4.16). In contrast, the encoding was much stronger in M1TJ and

S1TJ. This suggests that ALM may control L indirectly over longer timescales via a signal to

other brain regions, or is simply permissive for control of L. We consider the latter possibility

unlikely because stimulation of ALM drives licking (Komiyama et al., 2010; Li et al., 2015;
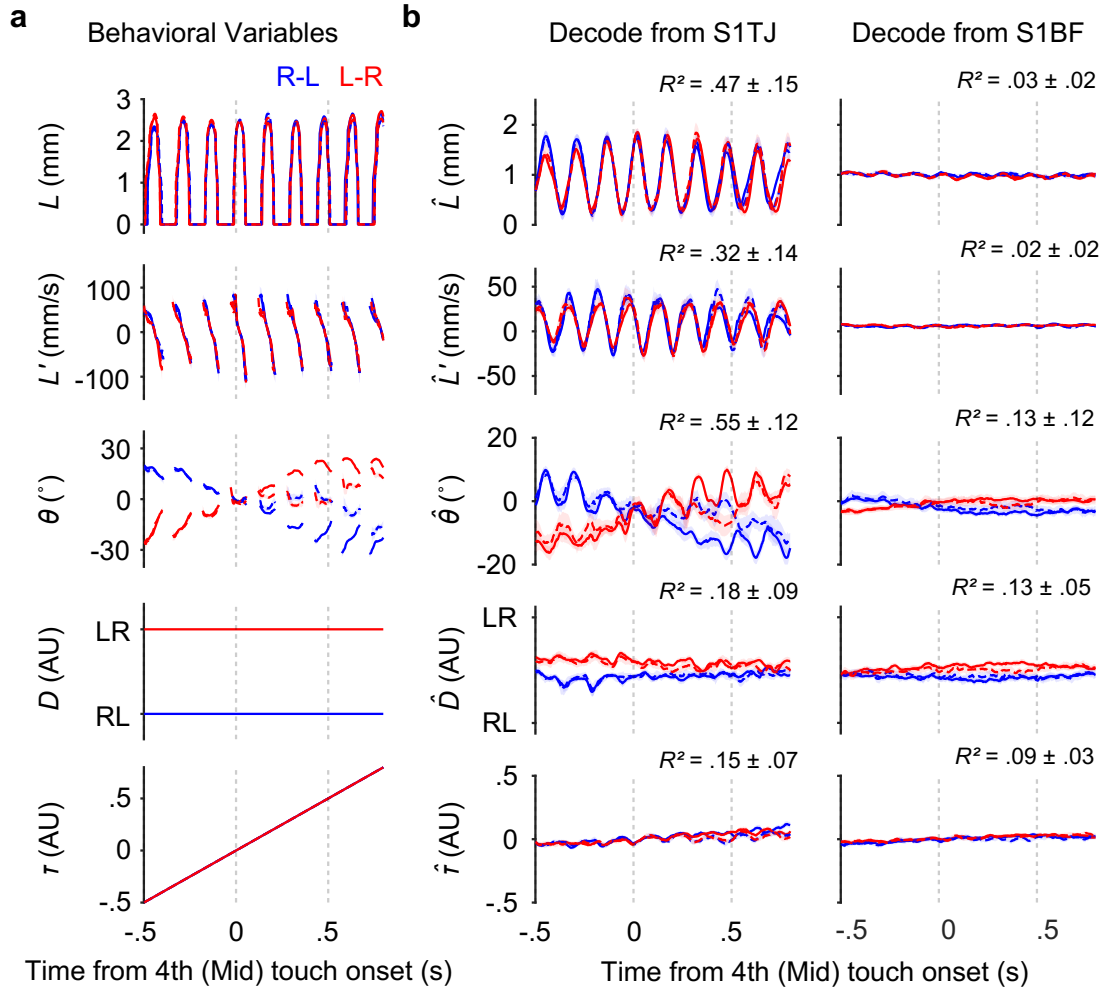
Figure 4.15: S1TJ and S1BF population decoding of behavioral variables

(a) The five recorded (or derived) behavioral variables averaged across trials (n = 2684 trials; mean ± 99% bootstrap confidence interval) with standard sequences (solid lines) or backtracking sequences (dashed lines), each of which is either right to left (blue) or left to right (red). Time along the x-axis is aligned to the fourth lick in the sequence. Time points where more than 80% of trials did not have observations are not plotted.

(b) Decoding of the five behavioral variables (rows) from populations recorded in S1TJ and S1BF (columns). Cross-validated R^2 for each region and variable is given (mean ± SD; S1TJ, n = 8 recordings; S1BF, n = 5 recordings). Same plotting conventions as in (**a**).

Allen et al., 2017). Our previous results showed that the pattern of L and L' of individual licks are the same regardless of sequence direction or whether or not mice were relocating a port during backtracking sequences (Figure 2.4a). The overlapping decoded L and L' traces for M1TJ and S1TJ are consistent with this behavioral invariance.
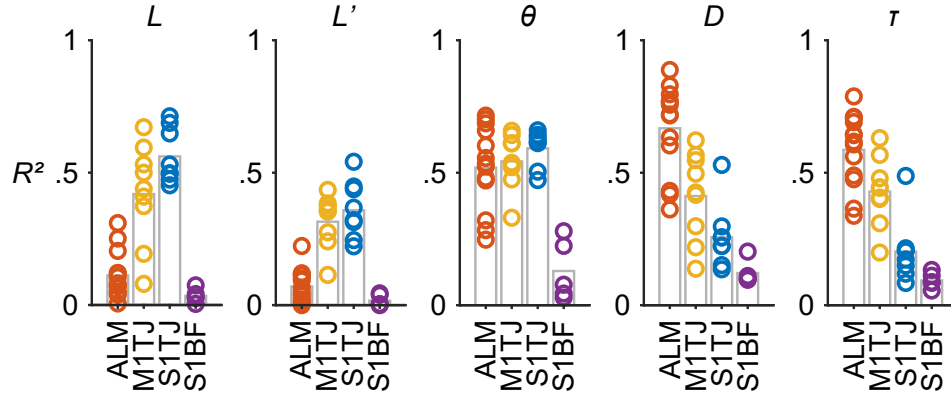
Figure 4.16: Goodness of fit across cortical regions

Goodness of fit for linear models that predict each of the five behavioral variables, quantified by cross-validated R^2. Each plot symbol shows one recording session.

ALM, M1TJ and S1TJ, but not S1BF, all showed strong encoding of θ (Figure 4.14, 4.15, and 4.16). When detecting backtracking, mice licked back to a previous angle to relocate the port and then progressed through the rest of the sequence. The opposing deflections in the decoded θ from backtracking trials (Figure 4.14 and 4.15; dashed curves for θ) matched this behavior. The traces of decoded θ in M1TJ and S1TJ contained rhythmic fluctuations that were absent in ALM, despite similar overall levels of encoding of θ (R^2 values). These fluctuations indicate that M1TJ and S1TJ encoded θ in a more instantaneous manner, whereas ALM encoded θ in a continuously modulated manner that may provide a control signal for the intended lick angle or represent the position of the target port.

Higher-level cortical regions are in part defined by the presence of more abstract (or latent) representations of sensory, motor and cognitive variables. Compared with L, L' and θ, which describe the kinematics of individual licks, sequence direction (D) and relative sequence time (τ) describe more abstract motor variables. In ALM we found the strongest encoding of both D and τ (Figure 4.14 and 4.16). Encoding of D and τ became progressively weaker in M1TJ, S1TJ, and S1BF, respectively. Overall, these results reveal a neural coding scheme with increasing levels of abstraction across S1TJ, M1TJ and ALM during the execution of flexible sensorimotor sequences.

4.3.3 Coding of Behavioral Variables for Sequence Generation Dominates Cortical Activity Patterns

Good decoding may come from a small fraction of informative units rather than dominant activity patterns across a population. More generally, we wondered whether the coding axes (i.e. coefficient vectors) of the five behavioral variables captured the dominant activity in the neural space, or were only related to minor components leaving the main dynamics unexplained. To determine this requires comparing the similarity between activity patterns captured by the coding axes and the dominant patterns in population activity identified in an unsupervised manner. In each recording session, we obtained neural trajectories in the coding subspaces (the subspaces spanned by coding axes) via linear decoding and trajectories in principal component (PC) subspaces (the subspaces spanned by the first few PCs) via PCA. Trajectories in PC subspaces depict dominant patterns in population activity but the PCs per se need not have any behavioral relevance. To see if neural trajectories in the coding and the PC subspaces were the same except for a change (rotation and/or scaling) in reference frame, we used canonical correlation (Methods) to find the linear transformation of the two trajectories such that they were maximally correlated.

After transformation, trajectories of the ALM population in the subspace of the top three PCs very well aligned (Figure 4.17a) and correlated (Figure 4.17b; group 2 in ALM) with the trajectories in the subspace encoding θ, D, and τ. This indicates that the most dominant neural activity patterns in the ALM population in fact encoded θ, D and τ. Since ALM minimally encoded L and L', including these to the coding subspaces decreased the correlation with PC trajectories (Figure 4.17b; group 1 and 3 in ALM). The decoded trajectories and PC trajectories in M1TJ and S1TJ also showed strong correlation but only when the coding subspaces included L and L'.

Across regions, the sum of variance explained (VE) by the five coding axes reached about half of that by the top five PCs (Methods; Figure 4.17d). The five coding axes were largely orthogonal among each other (Figure 4.17c). This indicates that they not only captured
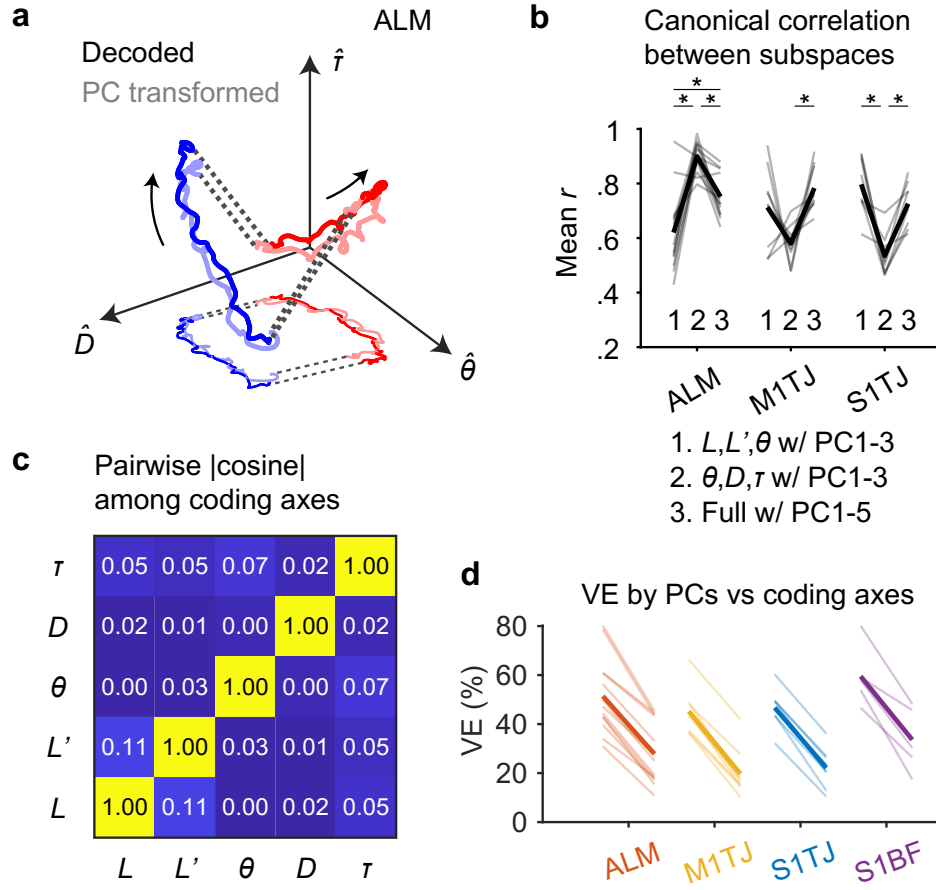
Figure 4.17: Coding of behavioral variables for sequence generation dominates cortical activity patterns

(a) Neural trajectories from ALM (mean; n = 13 recordings) during standard right to left (blue) and left to right (red) sequences (linked by dashed lines). Arrows indicate the direction of time. The decoded trajectories (darker thick curves) are overlaid with the trajectories (lighter thick curves) in the space of the top 3 PCs but linearly transformed to align with the decoded ones. A projection of these trajectories is shown in the D- plane with thinner and lighter curves.

(b) Mean canonical correlation coefficients (r) of each neural population (gray trace) across three conditions. The average mean r values for each condition are shown in black (ALM, n = 13 recordings; M1TJ, n = 9 recordings; S1TJ, n = 8 recordings). ∗ p < 0.001, not significant otherwise, paired two-tailed permutation test.

(c) Absolute pairwise cosine values among coding axes (mean; n = 35 recordings).

(d) Total percent variance explained (VE) by the first five principal components (left in each region) versus that by the five coding axes (right in each region) during sequence execution. Lighter lines show individual recording sessions and thicker lines show the means.

dominant neural dynamics but also did so efficiently with little redundancy.

4.4   Discussion

Individual neurons exhibit heterogeneous responses during sequence execution. We resisted to group neurons into categories based on artificially defined behavioral events (e.g. cue-responsive, left lick-preferred, or water-responsive neurons). There are many ways a neuron can respond to an event. For instance, a "cue-responsive" neuron can be transiently activated or inhibited, switched on or off, or modulated with a delay. Furthermore, it can be hard to deal with correlated behavioral events, such as auditory cue versus sequence initiation. Therefore, in order to reduce these interpretational biases, neuronal responses were algorithmically clustered in unsupervised ways based on the similarity among their PETHs without any behavioral labels. Then, we post-hoc interpret the functional implications of each cluster to the extent allowed by the task design.

Most of the neurons across ALM, M1TJ and S1TJ show selectivity to lick angle, many in conjunction to sequence direction. Neurons that are less discriminative respond at sequence initiation or throughout the period of execution. These responses not only temporally tile over the span of sequence execution but also tile differently in different sequences. The tiling is not an artifact that arises from random noise in the PETHs because 1) we excluded inactive neurons to avoid low signal-to-noise ratio responses and 2) the NNMF clustering was designed to capture patterns that best account for (or explain the greatest variance in) the population dynamics. Activities can tile a sequence with sparse code where neurons tend to fire precisely and exclusively to specific movements, or with dense code where responses tend to span several related movements at varying levels. In the high vocal center (HVC) of songbirds, neuronal activities tile a song with sharply and precisely fired spikes to specific syllables (Hahnloser et al., 2002). In contrast, the tiling in mice's licking sequences shows a much denser coding strategy which is similar to those observed in primate (e.g. Georgopoulos

et al. (1986)). The density of coding reflects a series of computational tradeoffs. Dense code allows the same number neurons to store more information and to be more robust against perturbations. However, representations in sparse code can be easily recombined to form new sequences (Ganguli and Sompolinsky, 2012; Bengio et al., 2013). The coding density may largely be constrained by the nature of the neural circuits, but can also be changed as a result of learning (Cao et al., 2015). It remains to be seen how the cortical representations change as mice gradually acquire the sequence behaviors through learning.

The seemingly continuous tiling from relatively heterogeneous neuronal responses resonates with a theory stating that most of the neurons are not made to directly control motor outputs, instead they work together in a dynamical system to support the unfolding of memorized patterns where only part of the patterns is transformed to output (Shenoy et al., 2013). Through decoding analyses, we reveal that neuronal populations in S1TJ, M1TJ and ALM encode behavioral variables with increasing levels of abstraction. The three variables (θ, D and τ) strongly encoded in ALM are not correlated with instantaneous kinematics of tongue motion, but reflect the need to keep track of the target position, lick angle update, and sequence progress. To uniquely represent a behavioral state at a moment within a given sequence, it is however sufficient to either encode θ or encode both D and τ. This apparent redundancy may serve a computational role from the dynamical systems perspective on motor control (Shenoy et al., 2013). Each point in the state space of a smooth dynamical system is associated with a unique trajectory (Strogatz, 2001). Thus, if the system needs to produce two licking sequences, the two trajectories must be kept adequately separated so that they do not intersect in the presence of noise and thereby turn one sequence into the other (Russo et al., 2018). Encoding all three variables avoids trajectories crossing at the middle lick and maximizes the distance of separation (Figure 4.17a).

Chapter 5

Reward Modulation and Cross-sequence Memory in ALM

5.1 Introduction

Sequence behaviors are composed of individual movements. It has been shown in the last Chapter (and reviewed in Chapter 1.2.1) that the brain contains not only signals to drive instantaneous movements but also signals about the context that determines proper sequencing. When multiple subsequences are organized into a higher level sequence, the state at the end of a subsequence must be persisted (and transformed) to correctly start the next subsequence. Most of the learned behavioral sequences are goal-oriented. Therefore, a natural way to inform the sequencing (though not always sufficient) is to keep track of a distance from the current behavioral state to the goal.

There are several formulations or related concepts for the distance to goal. The most normative account comes from reinforcement learning (RL) theories (Sutton and Barto, 2018). In this context, what measures the distance to goal is reward expectation, or value. An agent can learn a value function that maps the current behavioral state to its value by discounting the final reward given its delay in time and uncertainty. Once the reward is

achieved, there is no longer anything to expect and the reward expectation goes to zero. Such a value function can be learned via reward prediction errors signaled by the dopaminergic neurons in VTA and substantia nigra pars compacta (Watabe-Uchida et al., 2017) and acting upon the motor circuits to induce learning (Doya, 2000). Sometimes, the distance to goal can be simply reduced to the duration of time before obtaining it. For example, monkeys (Mita et al., 2009), rodents (Aronov et al., 2017) or artificial neural network-based agents (Deverett et al., 2019) have been trained in interval timing tasks where subjects needed to produce a delayed response by specific amounts of time. In these cases, a representation of time is sufficient to guide the motor production. However, although the time representation is accounted for by the temporal discounting in most RL algorithms, it should not be affected by a change in reward uncertainty.

Tasks for interval timing are closely related to motor preparation. The mouse ALM has recently become a focus for studying the circuit mechanisms of motor preparatory activity that link sensory cues to specific delayed motor responses (reviewed in Chapter 1.3.3). In a task where a stimulus and a response were separated by a delay, ALM neurons exhibited ramping activities encoding the future lick direction during the delay period (Guo et al., 2014; Li et al., 2016; Inagaki et al., 2019). Although preparatory activity in motor preparation has been relatively well-studied, it is not known whether similar mechanisms are playing a role in high level sequence organization where the information at the end of one subsequence persists to and instructs the next.

5.2 Results

5.2.1 Reward Modulates Dominant Activity Patterns in ALM

The coding axis of τ was identified by fitting models between neural activity and relative sequence time. However, if τ faithfully represents time, the downward deflection of traces from backtracking sequences (Figure 4.14b) should not appear, as time goes on regardless
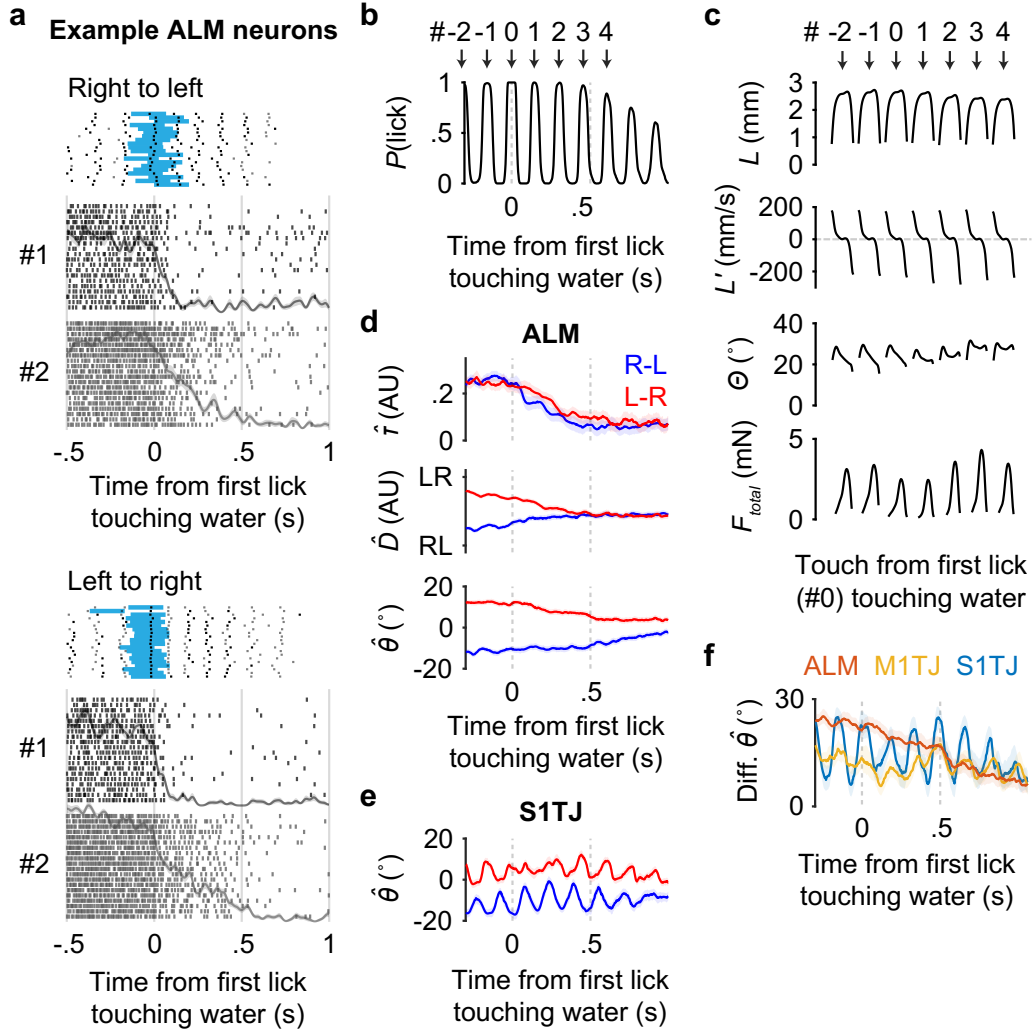
Figure 5.1: Reward modulation of activity in ALM

of what the animals do. Therefore, the patterns suggest a representation of a "distance to goal." In the context of the motor sequence, does the goal represent the initial arrival at the last port position, or the delivery of the water droplet, or finishing water consumption, etc?

In ALM, we found single neurons (Figure 5.1a) that fired actively during sequence execution but abruptly decreased their firing at the time when the tongue touched water droplets. Mice continued to emit ∼5 consummatory licks (Figure 5.1b) with similar, if not more strongly modulated, kinematics and lick force (Figure 5.1c). We decoded behavioral variables around the consumption period using linear models fitted during sequence execution (i.e. without data from the consumption period). The τ decoded from ALM populations

Figure 5.1: Reward modulation of activity in ALM

(a) Responses of two simultaneously recorded ALM neurons (#1 and #2) aligned at the first lick (specifically the middle of a tongue-out period) that touched water reward. For each sequence direction, shown at top are rasters of lick times (touches in black and misses in gray) and the duration of water delivery (blue) from 20 selected trials (Methods). Stacked below are spike rasters and the corresponding PETHs from the same 20 trials for each example neuron.

(b) The probability of licking (i.e. tongue-out) as a function of time. Licks are sequentially indexed with respect to the first lick (#0) touching the water.

(c) Patterns of kinematics and force for single licks around the first lick (#0) touching water (n = 25289 trials; mean ± 95% bootstrap confidence interval). The duration of individual licks was normalized. The total force (F_{total}) is the vector sum of vertical and lateral forces.

(d) Decoding of τ, D and θ from neuronal populations recorded in ALM (mean ± 99% bootstrap confidence interval; n = 13 recordings) in right to left (blue) or left to right (red) trials around the consumption period.

(e) Similar to (**d**) but decoding θ from S1TJ (n = 8 recordings).

(f) The difference between the decoded θ traces in right to left versus left to right trials (mean ± 99% bootstrap confidence interval; ALM, n = 13 recordings; M1TJ, n = 9 recordings; S1TJ, n = 8 recordings).

(Figure 5.1d; top) immediately decreased around the first contact (∼0 s) with water and diminished (∼0.5 s) before a noticeable decrease in lick production. This suggests that τ carried a reward expectation signal that smoothly increased as mice approached water delivery regardless of sequence direction or lick angle, but was suppressed by a delay of progress when backtracking occurred, and terminated at the time when the tongue detected water, despite continued licking movements. Curiously, the D coding (Figure 5.1d; middle) followed a similar time course as τ, although the implication of an interaction between sequence direction and reward is unclear.

The θ coding in ALM during the consumption period exhibited a more complex time course (Figure 5.1d, bottom). A moderate reduction in the separation of decoded θ in the two sequence directions occurred once the tongue touched the water. However, it remained separated toward the end of the lick bout. In contrast, the overall separation in θ coding from S1TJ (Figure 5.1e,f) and M1TJ (Figure 5.1f) was not altered by the detection of water, and the amplitude of rhythmic fluctuations was consistent with individual patterns of licking

force (compare S1TJ and M1TJ traces in Figure 5.1f with the F_{total} in Figure 5.1c). Overall, our results show that reward-related modulation is a prevalent feature in ALM but not S1TJ or M1TJ coding.

5.2.2 ALM Encodes Upcoming Sequences

Sequences with opposite directions were performed alternatively across trials (Figure 5.2a). Mice performed the task in darkness. Hairs and whiskers around the mouth were trimmed short to avoid contact with the port. Therefore, there was no external cue indicating which side a mouse should start a sequence. Nevertheless, expert mice were usually able to initiate sequences from the correct side without exploring the other (Figure 5.2b). This suggests that the information about target position (TP) was maintained internally during ITIs. Brain regions maintaining such information may be important in organizing individual sequences
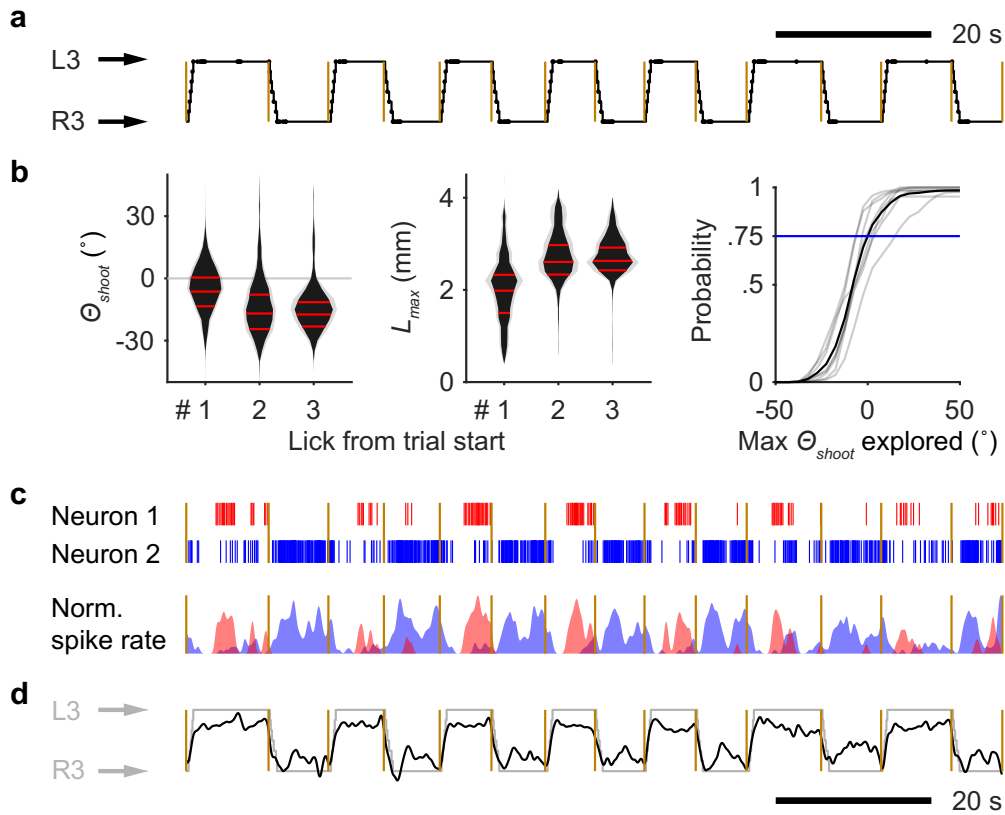
Figure 5.2: **ALM neurons encode sequences across trials**

Figure 5.2: ALM neurons encode sequences across trials

(a) Depiction of sequences performed by a mouse in alternating directions across 14 consecutive trials. Trial onsets are marked by yellow lines. Port positions shown in the black trace are overlaid with touch onsets (dots).

(b) Probability distributions of θ_{shoot} (left) and Lmax (right) for the first 3 licks at the start of a sequence (n = 8 mice; mean $\pm$ SD). The negative y-axis points to the side at which the port is located. The CDF (8 individual mice in gray and the mean in black) of the maximal θ_{shoot} explored before touching the port (at the side of negative θ). The blue line shows the probability of successfully locating the port without exploring beyond the midline.

(c) Top, rasters of two example units which had persistent and target position (TP) selective firing during the 14 consecutive trials in (**a**). Bottom, normalized and smoothed (0.25 s SD Gaussian kernel) spike rates of the two units.

(d) Decoded instantaneous TP (dark trace) from 58 simultaneously recorded units in ALM, overlaid with normalized port position (light trace).

together to form a higher level sequence.

In ALM recordings, we found simultaneously recorded units that fired persistently to specific TP during ITI (Figure 5.2c). A linear model fitted using data within 1 s before the cue onset showed smooth population decoding of TP across the span of many trials (Figure 5.2d). On average, ALM populations showed stronger encoding of TP (Figure 5.3a,b) compared to other regions, although overall it was weaker than either the encoding of D or θ during sequence execution. Interestingly, none of the regions, including ALM, encoded time or a distance to trial start (Figure 5.3c), perhaps because our ITI contained an exponential portion (Methods) whose flat hazard function made time to trial start unpredictable (Inagaki et al., 2019).

At the beginning of a sequence, D and θ both matched the preceding TP. Given that D and θ were encoded independently (i.e. in orthogonal neuronal subspaces) (Figure 4.17c), we wondered whether the coding of TP was more relevant to later D or θ, or neither. We used linear models fitted during the ITI to decode from neural activity during sequence execution. The resulting traces from two sequence directions crossed at the mid-sequence (Figure 5.3d). This revealed that TP coding reflects a persistent encoding of intended tongue angle in the absence of licks.
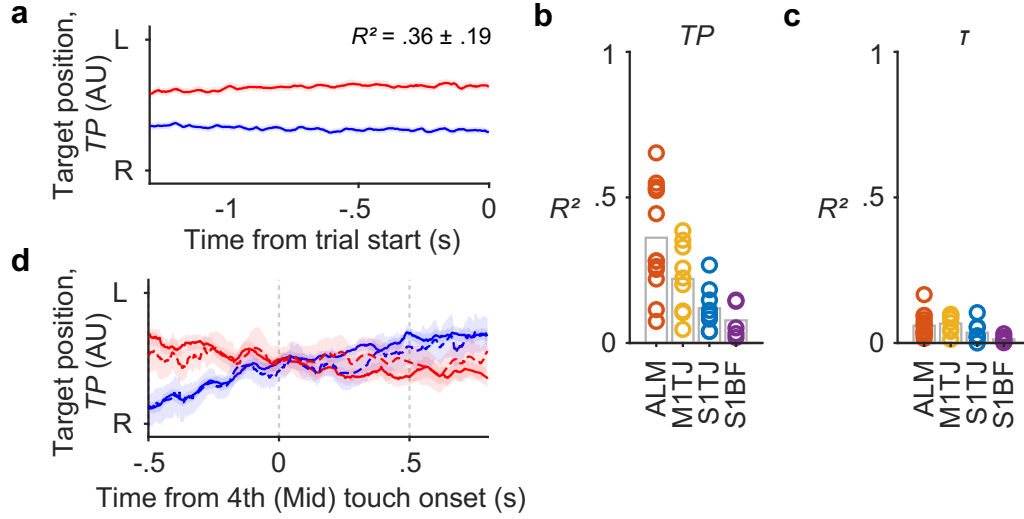
Figure 5.3: Coding of upcoming sequences in ALM

(a) Decoding of TP from ALM (mean $\pm$ 99% bootstrap confidence interval) before upcoming right to left trials (blue) or left to right trials (red). Cross-validated R^2 is shown (mean $\pm$ SD; n = 13 sessions).

(b) Goodness of fit for linear models that predict TP during ITIs, quantified by cross-validated R^2.

(c) Same as **(b)** but for τ.

(d) Using the same linear models in (Figure 5.2d) to decode TP during sequence execution. The decoded TP was averaged across trials (mean $\pm$ 99% bootstrap confidence interval) with standard sequences (solid lines) or backtracking sequences (dashed lines), each of which is either right to left (blue) or left to right (red).

5.3 Discussion

Premotor regions in the cerebral cortex are known to encode latent variables that are not directly correlated with specific motor outputs (Chapter 1.2.1 and 4). However, there has not been a unified view on the nature of these signals and how they come to exist.

The three main components of neural dynamics in ALM are target angle (θ), sequence direction (D) and sequence progress (τ) which are all subject to reward-related modulation, especially for τ. The changes in their activities are largely or partly independent from instantaneous licking kinematics and forces. The reward expectation-like signals in τ are consistent with a recent finding that ALM neurons ramp up activity as mice run towards a

105

reward zone (Chabrol et al., 2019), and are compatible with the ramping activities of ALM neurons observed in working memory tasks (Inagaki et al., 2019). Linking latent activities to an interplay between motor and reward representations is particularly attractive. It not only accounts for the phenomenology of observed activity patterns but also provides explanations as to how they can be acquired through learning.

To establish this link in the sequence licking behavior, additional experimental evidence and guidance from theories are needed. Although we show that single-neuron spiking and the population coding of τ drop at the water delivery, it is not clear whether this drop is triggered externally by the sensation of water or is generated internally due to the end of an expected reward time. To tie this loose end, one may look at trials with omitted water reward and see when the activity stops. In addition to tying loose ends, new insights may come from observing how neural representations evolve during learning and how the changes can be related to rewards.

Neurons in the mouse medial prefrontal cortex (mPFC) have been shown to encode working memory (Liu et al., 2014) and carry action values across trials (Bari et al., 2019). Here, we show that neurons in ALM bridge subsequences during ITIs by transforming last sequence direction to the lick angle of an upcoming sequence. It is not known whether mPFC is also involved in organizing sensorimotor sequences and what is its relationship with ALM. Functional brain imaging and perturbation experiments in humans suggest that the rostrolateral PFC mediates higher level sequence structure whereas the premotor cortex is only involved in subsequences (Desrochers et al., 2015b). Perhaps in mice, the functional cortical network may go beyond ALM to other prefrontal regions which mediate contextual information at longer timescales and with higher levels of abstraction.

Chapter 6

Conclusions and Perspectives

6.1 Main Conclusions

This thesis aims to understand how the brain generates and controls sensorimotor sequences. We show that mice can learn to perform flexible feedback- and memory-guided movements composed of individual licks that are organized into distinct sequences. Mice used their tongue to sequentially reach a set of predefined targets with remarkable speed and accuracy. When the target "backtracked" to a prior position unexpectedly, mice used tactile feedback to modify the ongoing sequence and quickly relocate the target. Mice alternated between two sequences across trials in the absence of external cues. This type of flexible feedback control, although common in daily life, differs from typical cerebellum based sensorimotor adaptation, where movements are fine tuned based on sensory prediction errors (Bastian, 2006; Shadmehr et al., 2010). The necessity of sensory feedback for ongoing execution also distinguishes our task from those involving repetitive or non-dexterous movement sequences (Jin and Costa, 2010; Kawai et al., 2015).

Closed-loop optogenetic inhibition and population single-unit electrophysiological recordings from multiple brain areas allowed us to identify three main cortical regions involved in controlling the sensorimotor sequences. The tongue premotor region (anterior lateral

motor cortex, ALM) was critical in initiating a sequence and controlling intended tongue angle smoothly across a series of individual licks. ALM neurons encoded sequence direction, progress toward the reward that signaled sequence completion, and intended tongue angle for upcoming sequences during the inter-trial interval, supporting a memory for how sequences were ordered within the behavioral session. In contrast, activity in the tongue/jaw region of the primary motor cortex (M1TJ) strongly encoded instantaneous lick angle, tongue length and tongue velocity. Tongue/jaw primary somatosensory cortex (S1TJ) activity encoded tongue angle, length and velocity, and its inhibition randomized lick angles, suggesting a proprioceptive role for these signals in S1TJ. Termination of licking depended on activity in multiple cortical regions.

Overall, our results reveal a functional sensorimotor cortical network that allows mice to use the tongue to perform complex and flexible motor sequences.

6.2   Sequence Licking Behavior in Mice

Eight years ago, I was in an international neuroscience meeting and heard a group of principal investigators with their wine glasses in hand ridiculing how mice were "blind" and how insensible it was to study vision in mice. Nowadays, studies of mouse vision have taken a large part of such meetings (perhaps even a bit overrepresented). The utility of a model organism is a multiplication between the biological functions the species exhibits and the experimental tools available for it. In the case of mouse vision, new studies have yet challenged the notion that mice only have very primitive and crude visual processing capabilities compared with primates' (Busse, 2018). However, it was the breakthrough in the available tools such as genetics-assisted circuit dissection and two-photon functional imaging during behaviors that opened up new possibilities of studying the circuit mechanisms of visual processing (e.g. Goard and Dan (2009); Li et al. (2012); Smith et al. (2013)) and general brain functions based on vision (e.g. Keller et al. (2012); Fiser et al. (2016); Zmarz and Keller (2016);

Leinweber et al. (2017)). Most of the experimental tools that benefit the studies of mouse vision equally apply to studies for motor control. However, is motor control in rodents qualitatively comparable to that in primates, or are rodents fundamentally "crippled" by primate standards?

Here, I want to distinguish two concepts - dexterity and (for the lack of a better word) complexity. We consider dexterity in terms of the degrees of freedom of a movement and how much speed and accuracy has to be traded off (Franklin and Wolpert, 2011). From recent adaptations of several primate motor tasks to rodents (Guo et al., 2015; Panigrahi et al., 2015; Mathis et al., 2017; Galiñanes et al., 2018), we have to admit that the dexterity of limb, paw, or saccadic movements in rodents is inferior to the primate counterparts. However, rodents are highly tactile species and use rich orofacial behaviors to interact with the environment (Diamond et al., 2008). Therefore, we may need to reconsider what dexterity means for the control of whiskers, tongue or jaw movements. Complexity refers to the organizing structure and cognitive demands of a behavior. Primates are capable of planning future movements, organizing primitive movements into more complex compositions, flexibly reconfiguring existing motor programs, and abstractly generalizing one learned sensorimotor pattern to another. It was not clear whether rodents are able to achieve behaviors with such complexity.

Through the novel task developed in this thesis, we found that mice exhibited impressive dexterity in tongue movements (especially the speed) and a remarkable ability to learn complex licking sequences. The sequence structure within a trial and the alternation of different sequences allow investigation of additional dimensions of motor control beyond simple stimulus-response contingencies (Crochet et al., 2019) or repetitive movement sequences (Jin and Costa, 2010). Mice were able to incorporate arbitrarily defined new transitions into the existing repertoire of learned sequences, and to quickly decide whether, and how, to change an ongoing motor program based on tactile feedback. These feedback-driven alterations involved complete reversals of direction and reorganization of the motor sequence. This novel

sequence licking task reveals previously unknown potential of the sensorimotor capabilities in mice and, I hope, will launch various new research directions.

6.3   A Functional Hierarchy for Sequence Generation and Control

Ongoing performance in our task, even in well-trained mice, required activity in the motor cortex. To perform a sequence of directed licks, the nervous system must generate motor commands to lengthen the tongue and modulate its angle. Optogenetic inhibition of ALM impaired both the length of each lick and the ability to direct licks away from the midline. These results partly support recent work showing that inhibition of ALM impairs aspects of tongue kinematics in a cued-licking task (Bollu et al., 2019). Electrophysiological recordings in ALM revealed strong encoding of θ but surprisingly not L or L' (which can indirectly produce L). This suggests that ALM may control L indirectly over longer timescales via a signal to other cortical or subcortical regions. In contrast to the rhythmically fluctuating θ trajectories decoded from M1TJ and S1TJ populations, ALM populations encoded an intended θ that varied smoothly over time. Our data do not rule out that ALM may also encode an intended L, or target distance. In our task mice used similar lick lengths across different port positions and we thus cannot distinguish signals for intended L versus task engagement. A task requiring mice to lick at different distances may help to resolve this question.

Precise motor outputs in skilled behaviors typically rely on sensory feedback (Scott et al., 2015; Franklin and Wolpert, 2011). Optogenetic inhibition of S1TJ produced distinct motor deficits from those observed with ALM inhibition. With S1TJ inhibition, mice could direct licks laterally to the normal extent and with normal length. However, lick angles became randomized and mice could no longer perform the correct sequences. Interestingly, neurons in S1TJ not only encoded θ, but also L and L'. This suggests that the system either does

not need sensory feedback about tongue length to guide the licking sequence, presumably because the required lick lengths were always the same and could in principle be controlled in open-loop, or that the inhibited L and L' signals in S1TJ could be compensated by redundant information in other brain regions.

Performing different sequences in our task not only required precise control of single licks but also the appropriate sequencing of different licks. Our data reveal encoding and control of behavior at multiple levels of abstraction across cortical areas. Zebra finch birdsong, in which syllables are organized into sequences to form songs, has been a powerful model for dissecting neural control over motor sequences and has revealed a specialized circuit for hierarchical motor control (Fee and Scharff, 2010). In primates, control of movements is distributed across multiple primary and premotor areas defined cytoarchitecturally and functionally. Medial premotor areas including the supplementary motor area (SMA) and pre-SMA play a prominent role in controlling the sequential organization of movements, while M1 activity may also be involved in sequence organization but in general is more closely related to imminent and ongoing movements (Tanji and Shima, 1994; Shima and Tanji, 2000; Lu and Ashe, 2005). While different cortical areas show functional distinctions regarding the level of abstraction of behavioral variables encoded, multiple somatomotor cortical areas contain corticospinal neurons and thus have the potential to exert fairly direct control of muscles. In rodents, the organization of motor areas and their functional relationships are so far less clear. Our results show a functional organization of somatomotor areas that reflects encoding and control of task-related behavioral variables with increasing levels of abstraction, with S1TJ and M1TJ activity more tightly linked to instantaneous kinematics, and ALM to variables at the sequence level, specifically a smooth representation of intended lick angle (θ), sequence direction (D), and distance to reward (τ).

Inability to stop an ongoing movement sequence can be as devastating as the inability to initiate one. When we optogenetically inhibited M1B at water consumption, mice showed difficulty in terminating an ongoing lick bout. This continued licking was not due to failure

of the tongue to reach the port to retrieve water. It is also unlikely that inhibition prevented mice from sensing the water because: 1) cortical representations of water taste can be found in gustatory cortex (Stapleton, 2006; Accolla et al., 2007), which we did not inhibit; 2) inhibiting any region in the anterior dorsal cortex not just somatosensory cortex, induced a similar deficit in sequence termination; and 3) decreasing the stimulus light intensity by half in ALM and S1TJ resulted in stronger deficits, not weaker. These results suggest that multiple anterior cortices play a role in active termination of an ongoing sequence, perhaps via convergence of activity on a downstream target such as the basal ganglia circuits that are critical for both starting and stopping movement sequences (Jin and Costa, 2010).

Prior work showed that ALM holds working memory or motor preparation signals for upcoming lick direction in delayed response tasks, bridging the sensory-to-motor transformation (reviewed in Chapter 1.3.3). Consistent with these results, we first found that optogenetic inhibition of ALM at cue onset strongly impaired the initiation of licking sequences. Before the go cue, ALM populations persistently encoded the upcoming sequence throughout the ITI. Furthermore, this signal is closely related to θ, but not D, suggesting that it represents the target position rather than a target sequence direction, despite the two being correlated at sequence initiation. Unlike previous tasks, our task does not have an explicit sensory event indicating the side on which to initiate a delayed sequence. However, the last port touch in a trial could be considered to serve as the "sensory event" and the location of the first lick in the next trial as the "response". Thus, our findings generalize the forms of delayed sensory-to-motor transformation involving ALM and provide insights into how the brain solves a sequence of sequences.

6.4 Computational Modeling of Sequence Licking

A common way to understand a neural circuit is either explicitly or implicitly through a functional diagram where, for instance, region A activates region B that feedbacks on A

which in turn disinhibits C and etc. Diagrams with such distinct nodes and fixed topology are certainly unable to account for the dynamics in a heterogenous and plastic brain, but can at best superficially describe the snapshot of an operating state. Physical sciences in the past few centuries have had great successes in explaining nature using the rules of change rather than describing how things are at a particular moment. In neuroscience, a superior model should be one that specifies how the substrates of the nervous system change during learning.

Directly measuring changes in neural circuits can be hard in behaving animals, but one may infer the rules of change by combining steady-state measurements with theories of learning. In the present study, we found that sequence progress (τ) exhibited properties of a reward expectation signal, and D and θ also showed reward-related modulations. This is consistent with recent work showing reward-related activity in ALM that depends on output from the cerebellar dentate nucleus (Chabrol et al., 2019). Putting these findings into the context of reinforcement learning (RL) frameworks (see Chapter 5.1), one can ask what circuit architectures, behavioral objectives, and learning rules can give rise to these observations (Richards et al., 2019).

Recent progress in modeling behaviors using deep neural networks (DNN) have provided great insights in, for example, the generation and utilities of place- and grid-cell representations (Banino et al., 2018), a gradient of meta-learning across brain regions (Botvinick et al., 2019), and a distributional code for value in VTA neurons (Dabney et al., 2020). Modeling cognitive tasks does not require sophisticated simulations of task environments because the behavioral inputs and outputs can be highly symbolic and abstracted (e.g. Bari et al. (2019)). In contrast, a realistic task environment must be simulated along with DNNs if the research of interest also involves sensory processing and motor control. This may be partly why DNN modeling for sensorimotor control has received relatively less attention (but see Chapter 1.2.3 for a few examples). The sequence licking task gives us excellent opportunities to model the neural computation from one end to the other. It is simple in terms of the fewer

113

degrees of freedom for motor output and low-dimensional sensory feedback. Meanwhile, it is complex enough in terms of the behavioral features involved, including the hierarchical sequence structure, the need for a memory of latent states, and the flexible motor programming based on sensory feedback. The Appendix III describes a simulator developed for the sequence licking task, which can interface with an agent powered by various artificial intelligence algorithms. This establishes the foundation for future neural network modeling.

6.5   Error Detection and Motor Branching

One of the most interesting open questions is how mice switch from standard to backtracking sequences based on the missing tactile feedback. Mechanisms that allow sensory feedback to integrate with unfolding motor programs is a long-standing and active area of research (Evarts and Tanji, 1976; Heindorf et al., 2018; Pruszynski et al., 2011; Scott et al., 2015; Stavisky et al., 2017). We can break down this process into two steps. Firstly, the nervous system must detect the need for a change and signal an error, presumably by comparing expected sensory inputs with the actual sensory feedback. Secondly, the sensory error signal is used to change the motor outputs in specific ways.

Generally speaking, the second step is better understood than the first. It is well-known that the climbing fibers from inferior olive neurons send error signals to the Purkinje cells in the cerebellar cortex and cause adaptive changes in motor outputs to minimize the sensory errors. However, this kind of error-minimizing adaptation occurs slowly over repeated trials (Bastian, 2006; Shadmehr et al., 2010), and is qualitatively different from rapid and flexible switching among distinct paths of action (Telgen et al., 2014; Krakauer et al., 2019). The basal ganglia-thalamocortical loop is thought to play an important role in movement selection (Jueptner and Weiller, 1998; Gurney et al., 2001). Interestingly, basal ganglia neurons are activated more when subjects attend to sensory stimuli (Kropotov and Etlinger, 1999). This suggests possible integrations between sensory and motor signals in voluntary actions.

It remains a mystery regarding where and how the sensory error signals are computed in the first place. Psychological theories (Norman and Shallice, 1986; Botvinick and Plaut, 2004; Cooper and Shallice, 2006) have proposed a supervisory controller which monitors automatic or procedural behaviors, handles exceptions, and keeps track of progress toward a goal. A lapse in supervised control may lead to the automatic execution of non-desired action. Many human functional imaging and non-human primate neurophysiological studies have identified task-related error signals in anterior cingulate cortex (ACC) (Desrochers et al., 2016). The predictive coding theory (Rao and Ballard, 1999; Bastos et al., 2012) and the popularization of single-cell calcium imaging in mice have motivated the discovery of error-encoding neurons across various sensory cortical regions (Keller et al., 2012; Makino and Komiyama, 2015; Fiser et al., 2016; Zmarz and Keller, 2016; Leinweber et al., 2017; Schneider et al., 2018). Specific findings have been briefly reviewed in Chapter 1.2.2. We still do not know how these error signals emerge and whether these neurons are responsible for informing the motor centers to change actions. However, finding such neurons can provide "anchor points" for following studies to expand from. Once the problem can be reduced to sensory processing and sensorimotor transformation, we will be able to take advantage of existing conceptual and experimental frameworks and to link a wealth of past literature on these subjects.

6.6 Coda

Together, our results from behavior, population electrophysiology and optogenetics allowed us to define a functional network within mouse sensory and motor cortices that governs execution of flexible, feedback-driven sensorimotor sequences.

Appendix I: Satellites Behavioral Control System

Modern behavioral tasks require automated control and data acquisition to handle the increased complexity of task design, the precision of delivering stimulus and recording responses, and the increased scale of experiments. A behavioral system can be divided into three parts. First, a program stored in a computer or a microcontroller specifies the logic of a task. Second, instructions from the program control electrical circuits and devices to physically interact with subjects. Third, data of these interactions are recorded and stored in specific formats. In the early days, scientists had to build each of these from scratch and customly assemble a complete system. Some systems were designed in ways very specific to the tasks of interest, whereas others were designed with greater abstraction that reflected general operating principles of many different tasks. The latter evolved to become standardized behavioral control systems for the field to adopt and build upon. However, stronger standardization often comes with the cost of reduced flexibility and greater dependency on the developers' maintenance. Overly abstract system design also increases the users' learning cost. Therefore, the question is how a system can balance standardization and flexibility.

Although standardization and flexibility generally go against each other, they are not in a "zero-sum" tradeoff. Efficient standards can allow for more flexibility given the same level of complexity and abstraction. On the other hand, the popularization of hobbyist electronics and programming tools have significantly lowered the cost for even a non-professional to build

a system from scratch. Therefore, the need for using extra standards has been decreased.

The Satellites behavioral control system developed in this thesis attempts to find a new balance between standardization and flexibility. The most unique (perhaps paradoxical) feature of this system is that it is entirely optional for implementing a behavioral task; it just makes implementing a task a lot easier and more reliable. The backbone of this system is based on the open-source platform called Arduino, which has been well-supported by a large community for more than a decade and is still improving. The Satellites library only provides standardized ways to handle communications (wired or wireless), manage the time flow, execute probabilistic control, synchronize devices, and etc. The Satellite Shield expands the input and output capabilities of the microcontroller with integrated output amplifiers, contact detection circuits, and wireless connection interfaces. Lastly, the SatellitesViewer provides a general-purpose graphical user interface (GUI) for users to send commands to and record data from the hardware. The app is also able to manage a network of controllers and computers for large scale behavioral training.

7.1   Satellites Library for Arduino

The Arduino platform has been widely adopted from high school science projects to professional engineering laboratories. Programming Arduino is easy to learn but can be hard to perfect. The code can get complicated and error prone when a user needs to deal with complex timing control, communication, synchronization, or several of them at the same time. The Satellites C++ library provides a collection of simple-to-use methods under the Satellites class that can streamline these processes. Additional companion libraries were developed as add-ons for more specialized purposes. The following is an overview of different modules in the Satellites class as well as other companion classes. For more details please refer to the code examples and other documentation.

117

7.1.1 Communication Protocol

Simple Arduino programs for task control are often written in a hardcoded way where parameters (e.g. reward duration, trial probability, lickport coordinates) are stored in the flash and cannot be changed by user input when the program is executing. Allowing users or external devices to modify the program on the fly provides significantly more flexibility for task control and should be a basic feature for a general-purpose system.

The communication part of Satellites library implements a specific, yet common, messaging convention for sending and receiving interpretable strings. Each message received by a device needs to have a format in "CmdName,123,456,789,...". "CmdName" is a string of letters which indicates what command it is. "123", "456" and "789" are parameters associated with this command. There can be as many parameters as you need or none at all. Components of this message are separated by delimiters (commas "," in this example).

When a computer sends a command in this format to Arduino, the command can be read by a user defined reader function in which the parameters are unloaded (getIndex and getValue) and interpreted one by one based on the command name (getCmdName). This reader function can be called explicitly in users' code, or implicitly when the task is put on hold using specifically designed methods for time control.

7.1.2 Time Control

Time control is key for structuring events in a behavioral task. The simplest is a delay after one event and before the next. Other time control may depend on feedback from subjects' responses. A task can be halted indefinitely or up to a time limited until the subject triggers a specific input. Conversely, a delay can be renewed every time the subject triggers a specific input. The Satellites class has three methods for each of these cases - delay, delayUntil, and delayContinue. Different from typical custom time control code, these methods will repetitively call the reader function when the program is waiting. In this way, the program can concurrently handle incoming user commands when the task is on hold

but temporarily hold back those commands when executing time critical operations such as delivering stimulus.

7.1.3   Utilities and Companion Libraries

It is common for a behavioral task to have non-deterministic components. For example, one of three different stimuli can be randomly chosen given certain probabilities. Or, a reward is omitted in one out of ten times. In these cases, a user can use the "choose" method to do the random sampling. Another type of randomness involves drawing a value from a probability distribution (sometimes capped by minimal and/or maximal values). The Interval class is designed for such sampling.

Other libraries were developed for specific tasks rather than for general use. The Zaber motor library allows users to use an essential set of high level APIs to drive the linear stages (up to two axes per instance). It was developed before the company Zaber released their official library. The advantage is the simplicity of use, whereas the main disadvantage is the lack of compatibility to certain types of motor and potentially weaker reliability (though we have not yet tested how reliable the official library is). The ManyStepper and ManyVibMotor classes were developed for minimalist control of barebone stepper motors (without dedicated driver) and vibratory DC motors, respectively.

7.2   Satellite Shield

A microcontroller is only an integrated circuit (IC), or a piece of silicon chip, that provides input and output capabilities and is responsible for executing a program. It is analogous to our central nervous system. To make this chip functional, it has to be hooked up to other peripheral circuits and ICs for physical interfaces (e.g. headers, ports), power regulation, signal conditioning, circuit protection, and etc. This collection is usually put on a single printed circuit board (PCB) and is called a development board. Commonly found devel-

opment boards include Arduino Uno, MEGA, and Teensy series. To continue our analogy, a development board is like a complete nervous system - central plus peripheral. However, nothing meaningful can be done before an individual has muscles, skin, eyes and so on. Similarly, additional circuits and devices need to be connected to a development board for controlling motors, driving displays, receiving Wifi data, and etc. Many of these applications are commonly needed and have standardized circuit solutions. These circuits are made into modules that can be used as off-the-shelf add-ons to development boards. In the Arduino world, those modules, in the form of PCBs, are called shields because they are typically stacked on top of (thus "shielding") a development board. Following the same idea, the Satellite Shield was developed for the Teensy series development boards. The Satellite Shield addresses several general needs in behavioral control including driving high current actuators, detecting contact responses, and data communication.

7.2.1   Digital Output Amplification

Microcontrollers are designed to process signals with low voltage and low current. Therefore, the signals must be amplified in order to operate power hungry devices. The Satellite Board provides ten digital output amplification (sometimes called relay) channels that can be optionally connected to ten pins on Teensy via jumpers.

The output voltages are determined by the power supplies to the amplifiers (Darlington transistor arrays, ULN2801A). The ten channels are separated into two banks of five. Each bank can either be supplied by the 5V that also powers the development board or by a separate supply connected to the power ports (VS1 or VS2). The amplifiers can take up to 50V of voltage supply but in practice they should operate well below that limit (e.g. < 30V). The output current is determined by the demand of external circuits and the maximal output current, whichever is lower. Six of the ten channels (T3 to T5 and T20 to T23) support up to 1 A of current and the rest four up to 500 mA.

It is necessary to clarify a few issues about the amplified outputs. Imagine a light bulb

being powered by a battery (or any DC source) and we want to control it by adding a switch. The switch can be added either on the positive side or the negative side. An amplifier in the Satellites Board acts just like a switch. Importantly, this switch is located on the negative (ground) side of the device being powered. That is why each output port has the same plus sign which is always connected to the positive terminal of voltage supply, but differently labeled minus sign. The numbers next to the minus signs indicate which Teensy pin controls the switch. As a result, when the switch is off, the minus terminals are floating and the voltage can have arbitrary values. A common mistake is to assume they have zero volt (i.e. connected to ground).

7.2.2 Contact Detection

Contacts, such as button pressing, screen tapping, and licking, is a common form of interaction between a subject and a behavioral system. In contrast to mechanical buttons, non-mechanical contact transducers have the advantages of flexible form factors, ease of integration, and long lifespan. There are three popular solutions for non-mechanical contact detection used in behavioral systems. The first is conductive contact detection where a subject acts as a mechanical switch that opens (no touch) or closes (with touch) the detection circuit. The second is capacitive detection which is based on the changing capacitance as an object contacts the subjects (or vice versa). The third is via strain transducers that measure tiny deformation caused by contacts.

The Satellite Board offers two channels of conductive contact detection. The general principle of operation is similar to previous systems (Svoboda lab, HHMI Janelia Research Campus; Slotnick 2009). A subject is connected to a high potential and the conductive object being touched is connected to a signal input (L0 or L1). The power for this detection circuit can be separately supplied by an isolated source (VB+ and VB-; B stands for battery though any isolated source will work). Any mechanical switch has electrical bouncing issues when a contact is weak and unstable. To handle bouncing during loose touches, each channel

has a built-in variable RC filter with a time constant from 0 to 50 ms.

7.2.3   Communication Interfaces

In addition to the serial ports directly available from Teensy development board, the Satellite Board provides two interfaces for wireless communication. One can connect to an ESP8266 Wifi module or other modules with compatible pinout. The other can connect to a XBee serial interface board.

7.3   SatellitesViewer
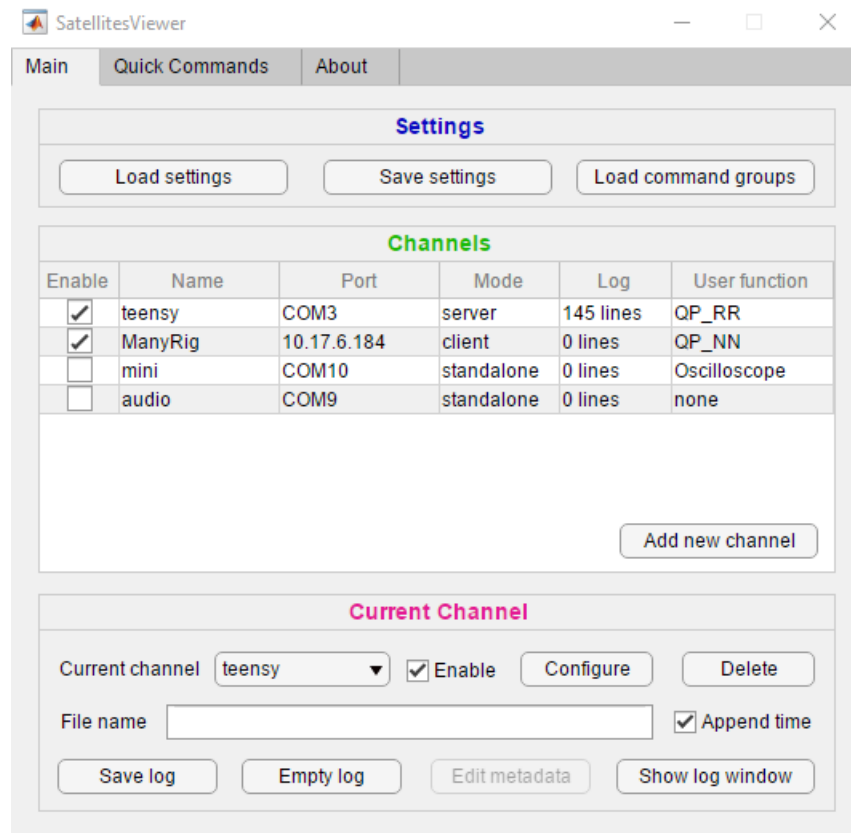
Figure 7.1: Main window of SatellitesViewer

Four channels were created in which two are currently enabled. The currently selected channel is "teensy".

SatellitesViewer is a general-purpose app for serial and UDP/IP communication. It

is developed using AppDesigner from MATLAB (MathWorks) and runs in the MATLAB environment. One can use it to simply communicate with a serial device (e.g. Arduino), or to talk with another computer over a local network, or to deploy multiple messaging topologies involving serial/Wifi/LAN interfaces. This app streamlines communications with easily customizable GUI (e.g. by editing an Excel spreadsheet). Its support of user functions makes online visualization and feedback control possible by basic MATLAB programming. The following explains the main concepts of operation.

7.3.1   Channel

SatellitesViewer uses the idea of Channel to manage a complete message flow based on a set of related connections. Users can create an arbitrary number of Channels. Each Channel has its independent configuration, input/output control, and logging, and can be accessed in the "Current Channel" panel. A Channel can be configured to one of three modes below.

Standalone mode

It is the simplest mode in which the computer connects to a device via a serial port (just like the Serial Monitor in Arduino IDE). However, one can, for example, create multiple Channels in this mode to control separate serial devices simultaneously. See "XBee API" section for advanced use of serial communication.

Server mode

This mode sets up a serial connection as in Standalone mode but also broadcasts messages received from this serial connection to one or more client computers in the local network (via UDP/IP). In the opposite direction, messages received from client computers will be relayed back to the serial device. One needs to provide the IP addresses and ports of client computers to send messages out and the port of the local computer to receive messages.
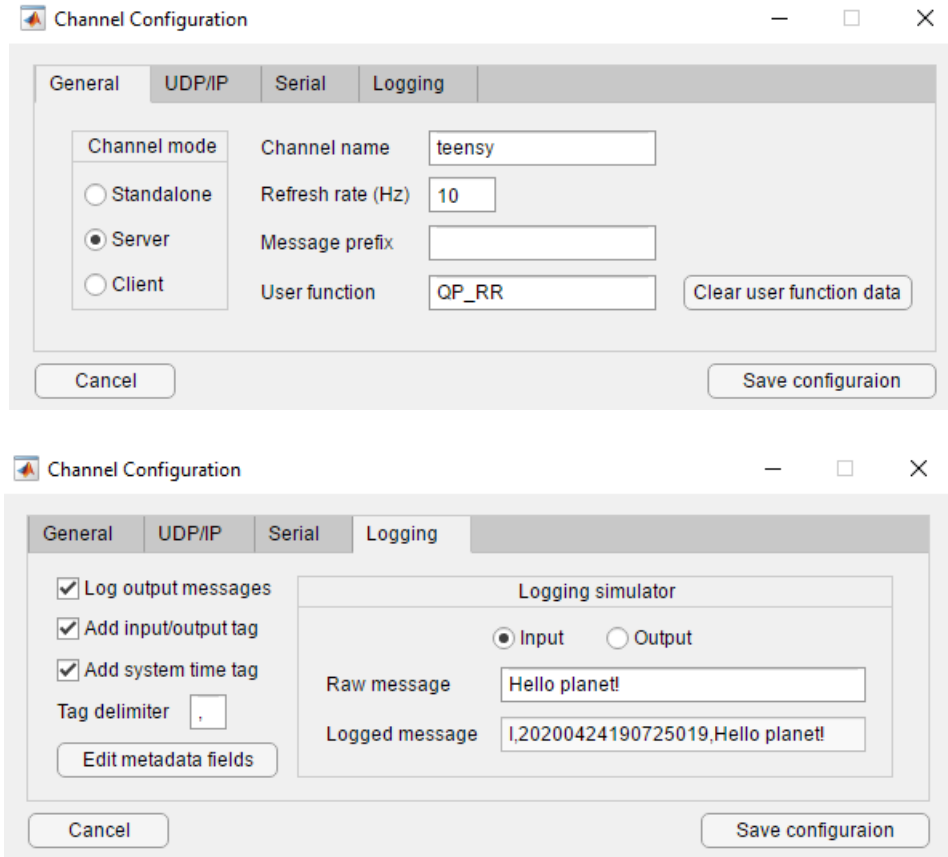
Figure 7.2: Adding or configuring a channel

Client mode

No serial connection is involved in this mode. Instead, the Channel sends and receives messages with a server computer on the network, thus can remotely control a serial device connected to the server. One needs to specify the IP address and port of the server computer to send messages out and the port of the local computer to receive messages. See "Client-Client communication" section for an alternative use.

Turn Channel on/off
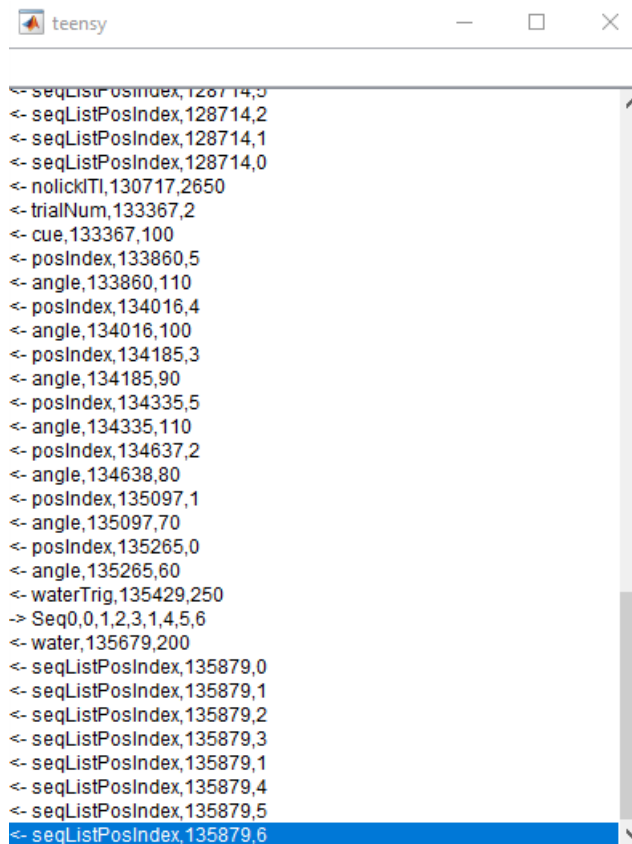
In the "Current Channel" panel, checking the Enable checkbox turns the current Channel on; unchecking turns it off. When turned off, a Channel has no actual connection and no computing resource is taken. See "Performance note" for more details on app performance

124

in different conditions.

7.3.2 Receiving Messages

Logging

Each Channel has its own communication log and a dedicated window for display. Operations such as saving and emptying logs can be accessed in the "Current Channel" panel. The saved log is a text file where each line is a message. More options about logging can be found in a Channel's configuration window. A simulator allows users to see how different options preprocess a message string.

Figure 7.3: Log window of a channel

The incoming messages are preceded by "<-" and the outgping commands by "->".

User function

One can write custom MATLAB functions to process incoming messages and achieve capabilities such as online visualization, closed-loop control, or even email notification. A user function takes one input argument and is called every time the Channel receives a message. This input argument is an object of type UserFuncModel which contains all available information and functionalities at your disposal. Please see example user functions in the "private" folder and documentation of UserFuncModel class for more details.
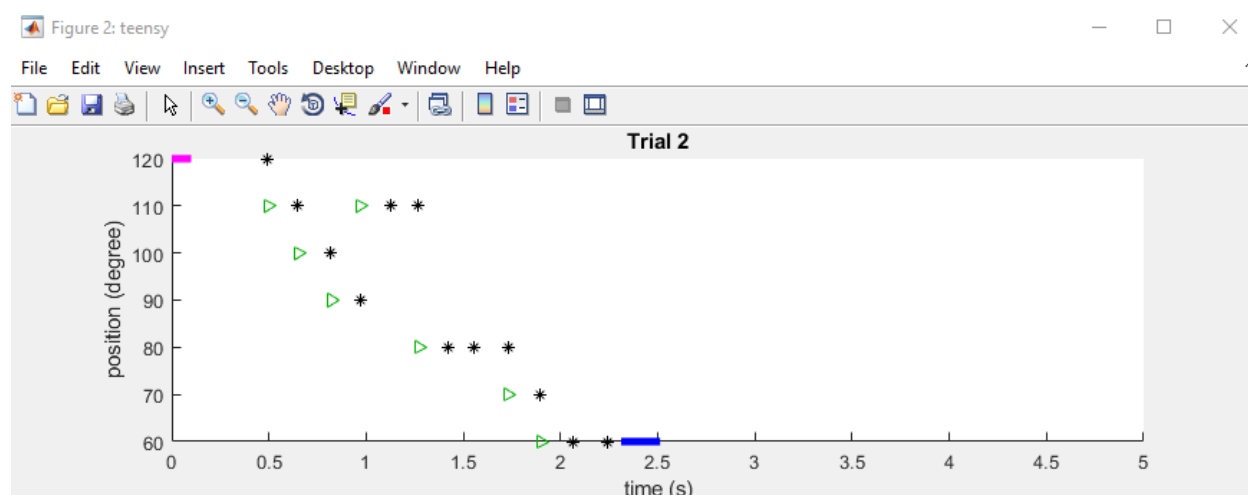
Figure 7.4: Online visualization of behavior using user function

This plot shows a trial in the sequence licking task. The magenta bar indicates the auditory cue. The blue bar indicates water reward delivery. Asterisks indicate lick onsets. Green triangles show when, and to where, the lickport was actuated upon each lick.

7.3.3 Sending Messages

Dialog

Messages can be sent in a Channel's log window just like that in Serial Monitor and many other messaging apps.

Quick commands

A user may have some frequently used commands to control a device and do not want to type them again and again. In the Quick Commands tab, one can input commonly used commands in individual boxes and send any of them by a click. Note that quick commands only send messages to the currently selected Channel.
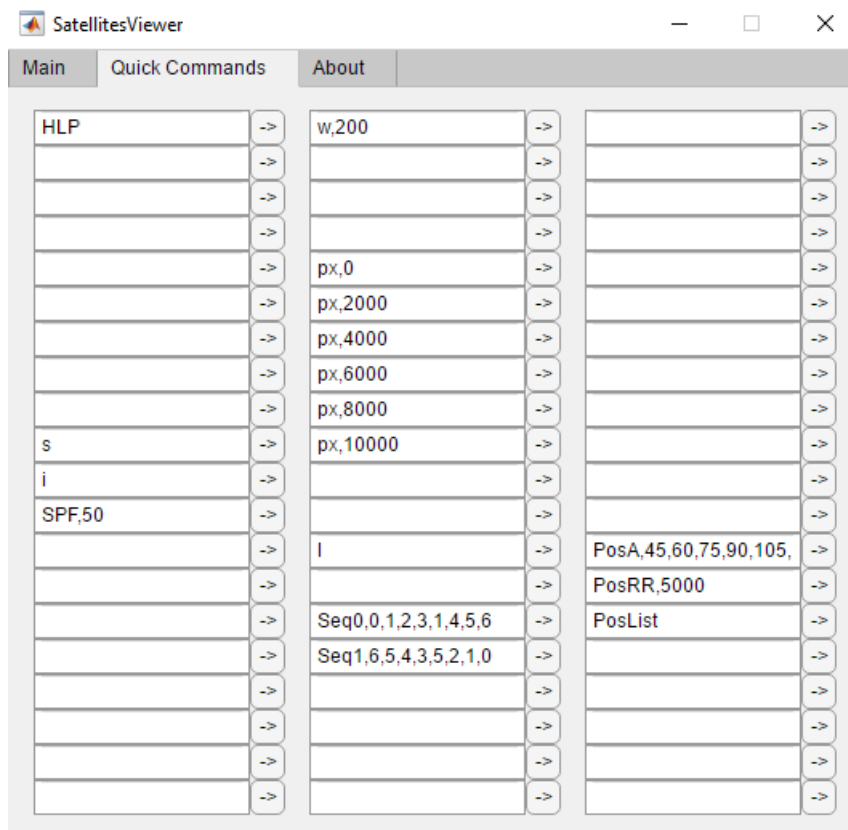
Figure 7.5: Quick Commands tab

Clicking a "->" button sends the command out.

Command groups

Quick commands are limited by the number of boxes, uniform layout, and lack of explanation. Command groups allow users to generate GUI with a large number of organized and well labeled custom commands. It converts content in Excel spreadsheets to buttons, textboxes and descriptive labels. Clicking a button sends a command and its values out using a format

specified by Satellites communication protocol. Playing with the example Excel file will help to understand the conventions. Like quick commands, command groups only send messages to the currently selected Channel.
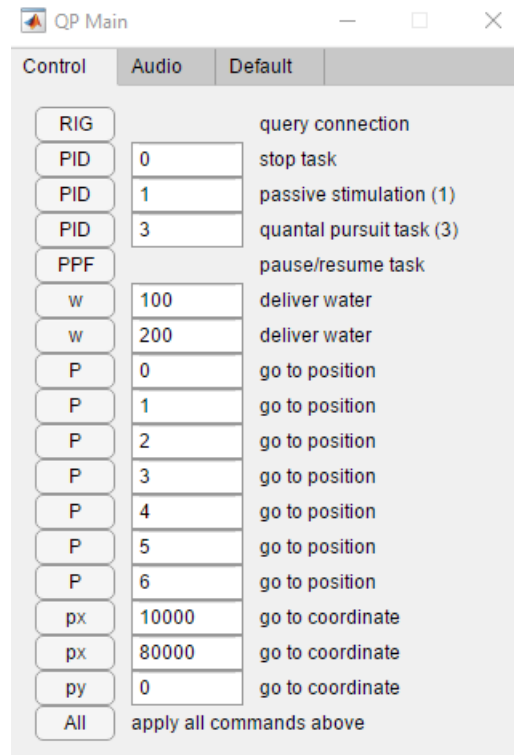
Figure 7.6: Example command groups

7.3.4 Save and Load Settings

Users can save all app settings to a MAT file using the "Save settings" button. This file can be loaded later to restore the state of the app. Please note that the following information is not considered as part of app settings: 1) messages in log, 2) log file name textbox and append time checkbox, 3) command groups, and 4) data in user functions.

7.3.5   Advanced Topics

XBee API mode

In addition to commonly used transparent serial communication, SatellitesViewer also supports a simplified version of XBee API mode allowing users to wirelessly communicate with one or more serial devices.

Appendix II: Optical Strain and Force Sensors

According to Newton's third law, the force exerted equals the force experienced. Being able to measure forces one exerts on an object not only tells us information carried by the motor command but also the resulting tactile feedback. Measuring forces directly is not easy. Most methods instead measure the strain (or deformation) in a material and then convert it into force based on known relationships. For example, a resistive force sensor relies on the change in a material's electrical resistance when deformed by external force. Piezoelectric sensors produce electric charges as the material deforms. Mechanical force gauges use a material's mechanical, instead of electrical, properties to relate strains to forces. The optical strain/force sensors developed in this thesis are in effect mechanical. The optical part does not mean it derives forces by measuring changes in a material's optical properties as a result of deformation. Instead, the optical components, photointerrupters, are used to measure displacement of the mechanical part with high spatial (submicron) and temporal (megahertz) resolution.

8.1 Mechanical Design

The core of this system is an optical sensor block. There are three mutually perpendicular channels inside the block. Two channels house photointerrupters (GP1S094HCZ0F, Sharp)

whose position can be manually adjusted by turning their corresponding rail screws. The third channel allows an object, such as a rod, to pass through. The object normally blocks about half of the light received by a photointerrupter, keeping the output voltage in the middle of the measurement range. Forces that cause relative displacements between the object and the photointerrupter change the amount of light being blocked thus changing the output voltage. Each photointerrupter can detect displacements in one dimension. When the two photointerrupters are both aligned on the object, they can detect the object's displacements in two perpendicular directions.

For the lickport sensors used in this thesis, the stainless steel lick tube is the target object. It is fixed on one end to form a cantilever. Mice licked the other free end, producing a small displacement ($<$ 0.1 mm at the tip for 5 mN) of the tube. Two photointerrupters placed perpendicular to the tube (Figure) convert the vertical and horizontal components of displacement into voltage signals. Pushing the tip down caused the cantilever to block more light at the vertical sensor and thereby decreased the output voltage; conversely, less force applied at the tip resulted in increased voltage. For the horizontal sensor, pushing the tube to the left or right decreased or increased the voltage output, respectively.

The optical force sensors can have different structural configurations based on the desired dynamic range and constraints in form factors. For example, to detect the body movement of mice, we yoked the mouse tube onto a cantilever going through the same type of sensor block. However, the forces exerted by the whole body is much larger than the force of licking. To avoid signal saturation, the displacement of the cantilever must be a lot smaller given the same amount of force (the displacement-response relationship of photointerrupters is fixed). There are two easy ways to achieve this - increasing the rigidity of the cantilever and reducing the mechanical advantage (e.g. length of the lever arm from the mouse tube) of force application. One can do the opposite to reduce the dynamic range and increase the sensitivity. For example, it was possible to detect the bending force of a single hair or a gentle air blow against the cantilever.

So far, all the structural parts of the sensor blocks were made by 3D printers (Ultimaker 3+) using PLA plastics as the material. The printers and the material do not offer high accuracy, or low systematic bias (typically < 200 µm with 0.4 mm nozzle or < 100 µm with 0.25 mm nozzle). To obtain the desired dimensions, one may need to empirically find the systematic bias by iteratively tweaking the dimensions in design. Although quite inconvenient, this is an addressable issue and usually takes one-time (or once in a while) effort as long as the printer remains stable. A bigger problem however is using plastic materials such as PLA to make parts that hold the object going through the sensor block. When a plastic material is deformed beyond a certain limit, it does not fully restore to the original shape after the force is removed. This is a material property called plasticity, not surprisingly. As a result, applying large forces back and forth will produce hysteresis in output signals, making any fixed force-voltage relationship invalid. The best way to address this issue is to replace plastic with metallic materials. Alternatively, one can keep the possible forces within a range where plasticity is negligible (e.g. the lickport case). The sensor block is a static structure during operation and does not bear a changing force. Therefore, there is little concern making it with plastics.

8.2   Signal Conditioning and Calibration

The receiver side of a photointerrupter is a phototransistor. The amount of light received by the photointerrupter is proportional to the amount of current (I) that can flow through it. To convert the current to voltage, a resistor (R) is put in series with the phototransistor. The voltage across that resistor is $V_R = I \times R = U \times R/(R + R_T)$, where U is the supply voltage and R_T is the effective resistance of the phototransistor. Since both R and the R_T contribute to the total resistance, if R is too small, V_R will be close to zero even if all the light passes through; if R is too big, V_R will be quite high even if all the light is blocked. The ideal case is to choose an appropriate R so that the minimal V_R is close to zero and the

max V_R is close to the supply voltage U.
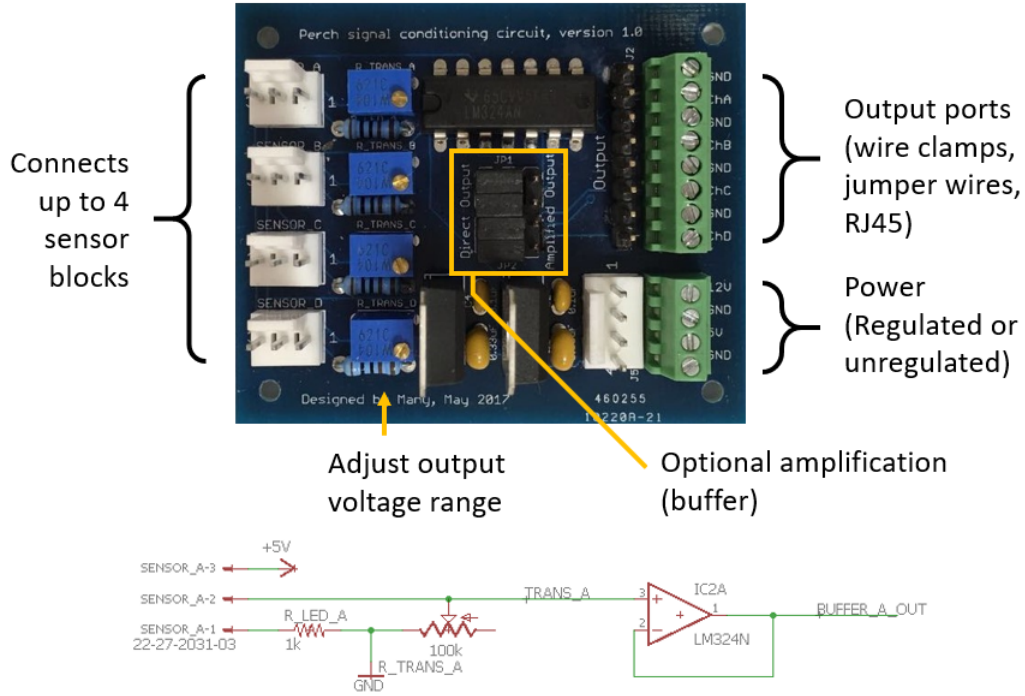
Figure 8.7: Perch signal conditioning board

On the signal conditioning board (Figure 8.7), R is a potentiometer with an adjustable resistance from 0 to 100kΩ. To find the best R, one first fully unblock the light and adjust the potentiometer so that the output signal V_R just starts to asymptote toward the supply voltage. Then fully block the light and adjust the potentiometer so that V_R starts to asymptote toward zero volt. The idea is to maximize the dynamic range of V_R. These adjustments must be made for every new photointerrupters due to their slightly variable electrical properties. If one finds it impossible to achieve a wide dynamic range or the output voltage is erratic, the phototransistor may have been damaged (e.g. due to excessive heat during soldering).

The signal conditioning board has four identical channels. Each channel can be set in one of two output modes, direct output or amplified output. Direct output directly outputs V_R, the voltage across R. However, this output has high output impedance and is prone to slow response time and ghosting. Amplified output first passes V_R to an op-amp used as an input buffer which then sends out signals with low output impedance.

133

In order to convert the output voltage to force, we need to experimentally calibrate their relationship. By design, the force and the output voltage follow a near linear relationship within a range of forces. To find this range, we measured the voltages (relative to baseline) with different weights added to the object. In the case of lickport sensors, this was achieved by hanging a centrifuge tube at the tip of the lickport and adding different amounts of water in it while recording the voltage response. Excellent linearity ($R^2 = 0.9999$) was achieved up to > 20 mN (Figure).

Appendix III: Sequence Licking Task Simulator

Research on reinforcement learning (RL) algorithms in artificial intelligence (AI) fields started by investigating how an algorithm performs in a specific and usually very simple task. However, real world RL agents including humans and mice are able to accomplish a variety of complex tasks using the same brains. To really say an RL algorithm has, or reflects, any "intelligence", one needs to benchmark it on different tasks and see if it is capable and versatile. One of the major breakthroughs in recent years was the development of DQN which shows human-level performance in many of the Atari 2600 games (Mnih et al., 2013, 2015). Aside from the successful integration of deep neural networks, one of the authors commented [1] (as I summarize) that the simulation of games has been key to the development as it offered a standardized benchmark directly comparable to human performance and covers a range of cognitive demands. Sharing similar spirits, the Psychlab was developed to simulate the laboratory environments in which human subjects or monkeys perform various vision-based tasks found in the literature of behavioral and cognitive neuroscience (Leibo et al., 2018). OpenAI gym is another platform that offers a standardized arena for simulating various control tasks (Brockman et al., 2016). Neuroscience research benefits from well-designed behavioral paradigms, and a collection of these have greatly advanced our understanding of the nervous system. Similarly, AI and computational neuroscience can

---

[1] Aug 19 2019. Episode 3: Life is like a game. *DeepMind: The Podcast.* `https://deepmind.com/blog/article/podcast-episode-3-life-is-like-a-game`

benefit by simulating and studying how artificial agents perform a collection of behavioral tasks (Botvinick et al., 2019; Yang et al., 2019).

The OpenAI gym focuses on simulating tasks that involve RL in episodic settings. As the name suggests, it allows the community to contribute tasks (or "environments") to the platform by implementing a standard set of application programming interfaces (API). An environment can be formalized as a partial observable Markov decision process (POMDP) where an agent makes an (usually incomplete) observation about the state of the environment, receives rewards or punishments if present, and in response outputs an action which may change the state of the environment. The incompleteness of observation may require an agent to produce different actions to the same observation via inference. This feature is present in the sequence licking task. For example, mice were able to remember where the lickport was left in the previous trial and initiate the next sequence from the correct sides although the auditory go-cue was always the same.

To simulate the sequence licking task, a Python class called QuantalPursuit was developed as an environment for OpenAI gym. (The name bears the history of the behavioral task development where mice were expected to follow, rather than drive, a moving lickport.) This environment implements the main interfaces such as *reset*, *step*, *observe*, *act*, *render*, and *close*. A custom Space class defines the specs of the observation vector and provides several utility methods including random sampling, normalization. The simulator has all the parameters used in the real task, such as port positions, port sequences, duration of water delay, cue duration, etc. The forces of touch feedback (F_{vert} and F_{lat}) are modeled to be proportional to the indentation (in mm) that the port goes into the tongue in each dimension. To implement the physical constraints of tongue motions, the tongue is not allowed to go beyond certain spatial limits (e.g. -60° and 60° for angle and 5.5 mm for length). The acceleration and deceleration of angle and length is also capped. Reward is controlled by four parameters. The first is the reward magnitude of water reward. The second is the cost of time. The other two impose cost when an agent extends the tongue long and wide,

respectively.

By changing the task parameters, one can make a simple licking task, a task developed in this thesis, or tasks with more complex sequence organization. As a proof of principle, here shows an agent with tabular Q-learning algorithm learned to perform a simple licking task in response to the cue (Figure 9.8).
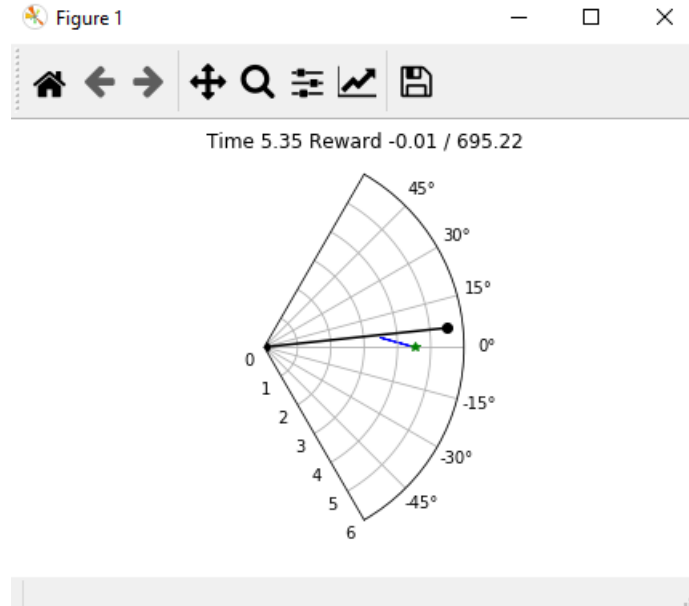
Figure 9.8: A snapshot of an performing agent

The agent is facing to the right, extending the "tongue" (black stem showing only its medial axis) out to reach the tip of the lickport (green star). The "tongue" has a width of 2 mm (not plotted) and is therefore touching the port. This produces a reaction force (blue arrow) pointing backward in the direction of contact and in proportion to the indentation.

Bibliography

Accolla R, Bathellier B, Petersen CCH, Carleton A. Differential Spatial Representation of Taste Modalities in the Rat Gustatory Cortex. *Journal of Neuroscience* 27: 1396–1404, 2007.

Agostino R, Berardelli A, Formica A, Accornero N, Manfredi M. SEQUENTIAL ARM MOVEMENTS IN PATIENTS WITH PARKINSON'S DISEASE, HUNTINGTON'S DISEASE AND DYSTONIA. *Brain* 115: 1481–1495, 1992. Publisher: Oxford Academic.

Allen WE, Kauvar IV, Chen MZ, Richman EB, Yang SJ, Chan K, Gradinaru V, Deverman BE, Luo L, Deisseroth K. Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex. *Neuron* 94: 891–907.e6, 2017.

Aoyagi H, Iwasaki Si, Nakamura K. Three-dimensional observation of mouse tongue muscles using micro-computed tomography. *Odontology* 103: 1–8, 2015.

Arce FI, Lee JC, Ross CF, Sessle BJ, Hatsopoulos NG. Directional information from neuronal ensembles in the primate orofacial sensorimotor cortex. *Journal of Neurophysiology* 110: 1357–1369, 2013.

Aronov D, Nevers R, Tank DW. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature* 543: 719–722, 2017. Number: 7647 Publisher: Nature Publishing Group.

Ayling OGS, Harrison TC, Boyd JD, Goroshkov A, Murphy TH. Automated light-based

mapping of motor cortex by photoactivation of channelrhodopsin-2 transgenic mice. *Nature Methods* 6: 219–224, 2009.

Badrinarayanan V, Kendall A, Cipolla R. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *arXiv:151100561 [cs]* , 2015. ArXiv: 1511.00561.

Banino A, Barry C, Uria B, Blundell C, Lillicrap T, Mirowski P, Pritzel A, Chadwick MJ, Degris T, Modayil J, Wayne G, Soyer H, Viola F, Zhang B, Goroshin R, Rabinowitz N, Pascanu R, Beattie C, Petersen S, Sadik A, Gaffney S, King H, Kavukcuoglu K, Hassabis D, Hadsell R, Kumaran D. Vector-based navigation using grid-like representations in artificial agents. *Nature* 557: 429–433, 2018.

Bari BA, Grossman CD, Lubin EE, Rajagopalan AE, Cressy JI, Cohen JY. Stable Representations of Decision Variables for Flexible Behavior. *Neuron* 103: 922–933.e7, 2019.

Barone P, Joseph JP. Prefrontal cortex and spatial sequencing in macaque monkey. *Experimental Brain Research* 78: 447–464, 1989.

Bastian AJ. Learning to predict the future: the cerebellum adapts feedforward movement control. *Current Opinion in Neurobiology* 16: 645–649, 2006.

Bastos A, Usrey W, Adams R, Mangun G, Fries P, Friston K. Canonical Microcircuits for Predictive Coding. *Neuron* 76: 695–711, 2012.

Benecke R, Rothwell JC, Dick JPR, Day BL, Marsden CD. DISTURBANCE OF SEQUENTIAL MOVEMENTS IN PATIENTS WITH PARKINSON'S DISEASE. *Brain* 110: 361–379, 1987. Publisher: Oxford Academic.

Bengio Y, Courville A, Vincent P. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35: 1798–1828, 2013. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.

Berdyyeva TK, Olson CR. Rank Signals in Four Areas of Macaque Frontal Cortex During Selection of Actions and Objects in Serial Order. *Journal of Neurophysiology* 104: 141–159, 2010. Publisher: American Physiological Society.

Berdyyeva TK, Olson CR. Intracortical microstimulation of supplementary eye field impairs ability of monkeys to make serially ordered saccades. *Journal of Neurophysiology* 111: 1529–1540, 2014. Publisher: American Physiological Society.

Bollu TP, Whitehead SC, Kardon B, Redd J, Liu MH, Goldberg JH. Tongue Kinematics. Cortex-dependent corrections as the mouse tongue reaches for, and misses, targets. *bioRxiv* p. 655852, 2019.

Bosma JF. *Development in the Fetus and Infant.* U.S. Department of Health, Education, and Welfare, National Institutes of Health, 1973.

Botvinick M, Plaut DC. Doing Without Schema Hierarchies: A Recurrent Connectionist Approach to Normal and Impaired Routine Sequential Action. *Psychological Review* 111: 395–429, 2004. Place: US Publisher: American Psychological Association.

Botvinick M, Ritter S, Wang JX, Kurth-Nelson Z, Blundell C, Hassabis D. Reinforcement Learning, Fast and Slow. *Trends in Cognitive Sciences* 23: 408–422, 2019.

Bowman JP, Aldes LD. Organization of the cerebellar tongue representation in the monkey. *Experimental Brain Research* 39: 249–259, 1980.

Brainard MS, Doupe AJ. Translating Birdsong: Songbirds as a Model for Basic and Applied Medical Research. *Annual Review of Neuroscience* 36: 489–517, 2013. _eprint: https://doi.org/10.1146/annurev-neuro-060909-152826.

Brockman G, Cheung V, Pettersson L, Schneider J, Schulman J, Tang J, Zaremba W. OpenAI Gym. *arXiv:160601540 [cs]* , 2016. ArXiv: 1606.01540.

Brown KJ. *Vocal Motor Coding in a Songbird Premotor Nucleus.* Ph.D., The University of Chicago, United States – Illinois, 2017. ISBN: 9780355519792.

Bryant JL, Boughter JD, Gong S, LeDoux MS, Heck DH. Cerebellar cortical output encodes temporal aspects of rhythmic licking movements and is necessary for normal licking frequency. *European Journal of Neuroscience* 32: 41–52, 2010. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1460-9568.2010.07244.x.

Busse L. Chapter 4 - The Mouse Visual System and Visual Perception. In Ennaceur A, de Souza Silva MA (eds.) *Handbook of Behavioral Neuroscience*, vol. 27 of *Handbook of Object Novelty Recognition*, pp. 53–68. Elsevier, 2018.

Cao V, Ye Y, Mastwal S, Ren M, Coon M, Liu Q, Costa R, Wang K. Motor Learning Consolidates Arc-Expressing Neuronal Ensembles in Secondary Motor Cortex. *Neuron* 86: 1385–1392, 2015.

Catalan MJ, Honda M, Weeks RA, Cohen LG, Hallett M. The functional neuroanatomy of simple and complex sequential finger movements: a PET study. *Brain* 121: 253–264, 1998. Publisher: Oxford Academic.

Chabrol FP, Blot A, Mrsic-Flogel TD. Cerebellar Contribution to Preparatory Activity in Motor Neocortex. *Neuron* 103: 506–519.e4, 2019.

Chen TW, Li N, Daie K, Svoboda K. A Map of Anticipatory Activity in Mouse Motor Cortex. *Neuron* 94: 866–879.e4, 2017.

Chesler AT, Szczot M, Bharucha-Goebel D, Čeko M, Donkervoort S, Laubacher C, Hayes LH, Alter K, Zampieri C, Stanley C, Innes AM, Mah JK, Grosmann CM, Bradley N, Nguyen D, Foley AR, Le Pichon CE, Bönnemann CG. The Role of PIEZO2 in Human Mechanosensation. *New England Journal of Medicine* 375: 1355–1364, 2016.

Clemens AM, Fernandez Delgado Y, Mehlman ML, Mishra P, Brecht M. Multisensory and Motor Representations in Rat Oral Somatosensory Cortex. *Scientific Reports* 8: 1–9, 2018. Number: 1 Publisher: Nature Publishing Group.

Clower WT. Early Contributions to the Reflex Chain Hypothesis. *Journal of the History of the Neurosciences* 7: 32–42, 1998.

Clower WT, Alexander GE. Movement Sequence-Related Activity Reflecting Numerical Order of Components in Supplementary and Presupplementary Motor Areas. *Journal of Neurophysiology* 80: 1562–1566, 1998. Publisher: American Physiological Society.

Cohen RG, Rosenbaum DA. Where grasps are made reveals how grasps are planned: generation and recall of motor plans. *Experimental Brain Research* 157: 486–495, 2004.

Cooper RP, Shallice T. Hierarchical schemas and goals in the control of sequential behavior. *Psychological Review* 113: 887–916, 2006. Place: US Publisher: American Psychological Association.

Crochet S, Lee SH, Petersen CC. Neural Circuits for Goal-Directed Sensorimotor Transformations. *Trends in Neurosciences* 42: 66–77, 2019.

Dabney W, Kurth-Nelson Z, Uchida N, Starkweather CK, Hassabis D, Munos R, Botvinick M. A distributional code for value in dopamine-based reinforcement learning. *Nature* 577: 671–675, 2020. Number: 7792 Publisher: Nature Publishing Group.

Davison AC, Hinkley DV. *Bootstrap methods and their application.* Cambridge ; New York, NY, USA: Cambridge University Press, 1997.

Delfs JM, Kelley AE. The role of D1 and D2 dopamine receptors in oral stereotypy induced by dopaminergic stimulation of the ventrolateral striatum. *Neuroscience* 39: 59–67, 1990.

Deschênes M, Takatoh J, Kurnikova A, Moore J, Demers M, Elbaz M, Furuta T, Wang F,

Kleinfeld D. Inhibition, Not Excitation, Drives Rhythmic Whisking. *Neuron* 90: 374–387, 2016.

Desrochers TM, Amemori Ki, Graybiel AM. Habit Learning by Naive Macaques Is Marked by Response Sharpening of Striatal Neurons Representing the Cost and Outcome of Acquired Action Sequences. *Neuron* 87: 853–868, 2015a. Publisher: Elsevier.

Desrochers TM, Burk DC, Badre D, Sheinberg DL. The Monitoring and Control of Task Sequences in Human and Non-Human Primates. *Frontiers in Systems Neuroscience* 9, 2016. Publisher: Frontiers.

Desrochers TM, Chatham CH, Badre D. The Necessity of Rostrolateral Prefrontal Cortex for Higher-Level Sequential Behavior. *Neuron* 87: 1357–1368, 2015b. Publisher: Elsevier.

Deuker L, Bellmund JL, Navarro Schröder T, Doeller CF. An event map of memory space in the hippocampus. *eLife* 5: e16534, 2016. Publisher: eLife Sciences Publications, Ltd.

Deverett B, Faulkner R, Fortunato M, Wayne G, Leibo JZ. Interval timing in deep reinforcement learning agents. *arXiv:190513469 [cs]* , 2019. ArXiv: 1905.13469.

Diamond ME, von Heimendahl M, Knutsen PM, Kleinfeld D, Ahissar E. 'Where' and 'what' in the whisker sensorimotor system. *Nature Reviews Neuroscience* 9: 601–612, 2008. Number: 8 Publisher: Nature Publishing Group.

Dick JPR, Benecke R, Rothwell JC, Day BL, Marsden CD. Simple and complex movements in a patient with infarction of the right supplementary motor area. *Movement Disorders* 1: 255–266, 1986. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mds.870010405.

Doran GA. Review of the evolution and phylogeny of the mammalian tongue. *Acta Anatomica* 91: 118–129, 1975.

Doya K. Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology* 10: 732–739, 2000.

Economo MN, Viswanathan S, Tasic B, Bas E, Winnubst J, Menon V, Graybuck LT, Nguyen TN, Smith KA, Yao Z, Wang L, Gerfen CR, Chandrashekar J, Zeng H, Looger LL, Svoboda K. Distinct descending motor cortex pathways and their roles in movement. *Nature* 563: 79–84, 2018. Number: 7729 Publisher: Nature Publishing Group.

Evarts EV, Tanji J. Reflex and intended responses in motor cortex pyramidal tract neurons of monkey. *Journal of Neurophysiology* 39: 1069–1080, 1976.

Fee MS, Scharff C. The songbird as a model for the generation and learning of complex sequential behaviors. *ILAR journal* 51: 362–377, 2010.

Felleman DJ, Van Essen DC. Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cerebral Cortex* 1: 1–47, 1991.

Fink AJP, Croce KR, Huang ZJ, Abbott LF, Jessell TM, Azim E. Presynaptic inhibition of spinal sensory feedback ensures smooth movement. *Nature* 509: 43–48, 2014. Number: 7498 Publisher: Nature Publishing Group.

Fiser A, Mahringer D, Oyibo HK, Petersen AV, Leinweber M, Keller GB. Experience-dependent spatial expectations in mouse visual cortex. *Nature Neuroscience* 19: 1658–1664, 2016.

Fonseca E, de Lafuente V, Simon SA, Gutierrez R. Sucrose intensity coding and decision-making in rat gustatory cortices. *eLife* 7: e41152, 2018. Publisher: eLife Sciences Publications, Ltd.

Foster DJ. Replay Comes of Age. *Annual Review of Neuroscience* 40: 581–602, 2017.

Franklin D, Wolpert D. Computational Mechanisms of Sensorimotor Control. *Neuron* 72: 425–442, 2011.

Fujii N, Graybiel AM. Representation of Action Sequence Boundaries by Macaque Prefrontal

Cortical Neurons. *Science* 301: 1246–1249, 2003. Publisher: American Association for the Advancement of Science Section: Report.

Gaffield MA, Christie JM. Movement Rate Is Encoded and Influenced by Widespread, Coherent Activity of Cerebellar Molecular Layer Interneurons. *Journal of Neuroscience* 37: 4751–4765, 2017. Publisher: Society for Neuroscience Section: Research Articles.

Galiñanes GL, Bonardi C, Huber D. Directional Reaching for Water as a Cortex-Dependent Behavioral Framework for Mice. *Cell Reports* 22: 2767–2783, 2018.

Gallivan JP, Chapman CS, Wolpert DM, Flanagan JR. Decision-making in sensorimotor control. *Nature Reviews Neuroscience* 19: 519–534, 2018.

Ganguli S, Sompolinsky H. Compressed Sensing, Sparsity, and Dimensionality in Neuronal Information Processing and Data Analysis. *Annual Review of Neuroscience* 35: 485–508, 2012. _eprint: https://doi.org/10.1146/annurev-neuro-062111-150410.

Gao Z, Davis C, Thomas AM, Economo MN, Abrego AM, Svoboda K, De Zeeuw CI, Li N. A cortico-cerebellar loop for motor planning. *Nature* 563: 113–116, 2018.

Gavornik JP, Bear MF. Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nature Neuroscience* 17: 732–737, 2014.

Georgopoulos A, Schwartz A, Kettner R. Neuronal population coding of movement direction. *Science* 233: 1416–1419, 1986.

Georgopoulos AP. Neural aspects of cognitive motor control. *Current Opinion in Neurobiology* 10: 238–241, 2000.

Goard M, Dan Y. Basal forebrain activation enhances cortical coding of natural scenes. *Nature Neuroscience* 12: 1444–1449, 2009. Number: 11 Publisher: Nature Publishing Group.

Goldberg G. Supplementary motor area structure and function: Review and hypotheses. *Behavioral and Brain Sciences* 8: 567–588, 1985. Publisher: Cambridge University Press.

Graham DM, Sun C, Hill DL. Temporal Signatures of Taste Quality Driven by Active Sensing. *Journal of Neuroscience* 34: 7398–7411, 2014. Publisher: Society for Neuroscience Section: Articles.

Grill HJ, Norgren R. The taste reactivity test. I. Mimetic responses to gustatory stimuli in neurologically normal rats. *Brain Research* 143: 263–279, 1978.

Grosbras MH, Leonards U, Lobel E, Poline JB, LeBihan D, Berthoz A. Human Cortical Networks for New and Familiar Sequences of Saccades. *Cerebral Cortex* 11: 936–945, 2001. Publisher: Oxford Academic.

Guo JZ, Graves AR, Guo WW, Zheng J, Lee A, Rodríguez-González J, Li N, Macklin JJ, Phillips JW, Mensh BD, Branson K, Hantman AW. Cortex commands the performance of skilled movement. *eLife* 4: e10774, 2015.

Guo Z, Li N, Huber D, Ophir E, Gutnisky D, Ting J, Feng G, Svoboda K. Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* 81: 179–194, 2014.

Guo ZV, Inagaki HK, Daie K, Druckmann S, Gerfen CR, Svoboda K. Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* 545: 181–186, 2017.

Gurney K, Prescott TJ, Redgrave P. A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological Cybernetics* 84: 401–410, 2001.

Hahnloser RHR, Kozhevnikov AA, Fee MS. An ultra-sparse code underliesthe generation of neural sequences in a songbird. *Nature* 419: 65–70, 2002.

Halsband U, Ito N, Tanji J, Freund HJ. The role of premotor cortex and the supplementary motor area in the temporal control of movement in man. *Brain* 116: 243–266, 1993. Publisher: Oxford Academic.

Harrington DL, Haaland KY. SEQUENCING IN PARKINSON'S DISEASEABNORMALITIES IN PROGRAMMING AND CONTROLLING MOVEMENT. *Brain* 114A: 99–115, 1991. Publisher: Oxford Academic.

Hayhoe M, Ballard D. Eye movements in natural behavior. *Trends in Cognitive Sciences* 9: 188–194, 2005.

Hayhoe MM. Vision and Action. *Annual Review of Vision Science* 3: 389–413, 2017. _eprint: https://doi.org/10.1146/annurev-vision-102016-061437.

He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. *arXiv:151203385 [cs]* , 2015. ArXiv: 1512.03385.

Heindorf M, Arber S, Keller GB. Mouse Motor Cortex Coordinates the Behavioral Response to Unpredicted Sensory Feedback. *Neuron* 99: 1040–1054.e5, 2018.

Hikosaka O, Nakamura K, Sakai K, Nakahara H. Central mechanisms of motor skill learning. *Current Opinion in Neurobiology* 12: 217–222, 2002.

Hikosaka O, Rand MK, Miyachi S, Miyashita K. Learning of sequential movements in the monkey: process of learning and retention of memory. *Journal of Neurophysiology* 74: 1652–1661, 1995. Publisher: American Physiological Society.

Hill DN, Mehta SB, Kleinfeld D. Quality Metrics to Accompany Spike Sorting of Extracellular Signals. *Journal of Neuroscience* 31: 8699–8705, 2011.

Hunnicutt BJ, Jongbloets BC, Birdsong WT, Gertz KJ, Zhong H, Mao T. A comprehensive excitatory input map of the striatum reveals novel functional organization. *eLife* 5: e19103, 2016. Publisher: eLife Sciences Publications, Ltd.

Inagaki HK, Fontolan L, Romani S, Svoboda K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* 566: 212–217, 2019.

Inchul P, Amano N, Satoda T, Murata T, Kawagishi S, Yoshino K, Tanaka K. Control of oro-facio-lingual movements by the substantia nigra pars reticulata: High-frequency electrical microstimulation and GABA microinjection findings in rats. *Neuroscience* 134: 677–689, 2005.

Isoda M, Tanji J. Cellular Activity in the Supplementary Eye Field During Sequential Performance of Multiple Saccades. *Journal of Neurophysiology* 88: 3541–3545, 2002. Publisher: American Physiological Society.

Jenkins IH, Brooks DJ, Nixon PD, Frackowiak RS, Passingham RE. Motor sequence learning: a study with positron emission tomography. *Journal of Neuroscience* 14: 3775–3790, 1994. Publisher: Society for Neuroscience Section: Articles.

Jiang X, Long T, Cao W, Li J, Dehaene S, Wang L. Production of Supra-regular Spatial Sequences by Macaque Monkeys. *Current Biology* 28: 1851–1859.e4, 2018.

Jin X, Costa RM. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature* 466: 457–462, 2010.

Jin X, Costa RM. Shaping action sequences in basal ganglia circuits. *Current Opinion in Neurobiology* 33: 188–196, 2015.

Jin X, Tecuapetla F, Costa RM. Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences. *Nature Neuroscience* 17: 423–430, 2014.

Jueptner M, Weiller C. A review of differences between basal ganglia and cerebellar control of movements as revealed by functional imaging studies. *Brain* 121: 1437–1449, 1998. Publisher: Oxford Academic.

Kami A, Meyer G, Jezzard P, Adams MM, Turner R, Ungerleider LG. Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature* 377: 155–158, 1995. Number: 6545 Publisher: Nature Publishing Group.

Kawai R, Markman T, Poddar R, Ko R, Fantana A, Dhawale A, Kampff A, Ölveczky B. Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* 86: 800–812, 2015.

Keller G, Bonhoeffer T, Hübener M. Sensorimotor Mismatch Signals in Primary Visual Cortex of the Behaving Mouse. *Neuron* 74: 809–815, 2012.

Kiehn O. Decoding the organization of spinal circuits that control locomotion. *Nature Reviews Neuroscience* 17: 224–238, 2016.

Kimura D. Left-hemisphere control of oral and brachial movements and their relation to communication. *Philosophical Transactions of the Royal Society of London B, Biological Sciences* 298: 135–149, 1982.

Kleinfeld D, Deschênes M, Moore JD. The Central Pattern Generator for Rhythmic Whisking. In Krieger P, Groh A (eds.) *Sensorimotor Integration in the Whisker System*, pp. 149–165. New York, NY: Springer, 2015.

Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Qi XL, Romo R, Uchida N, Machens CK. Demixed principal component analysis of neural population data. *eLife* 5: e10989, 2016. Publisher: eLife Sciences Publications, Ltd.

Koechlin E, Ody C, Kouneiher F. The Architecture of Cognitive Control in the Human Prefrontal Cortex. *Science* 302: 1181–1185, 2003. Publisher: American Association for the Advancement of Science Section: Research Article.

Kolb B, Milner B. Performance of complex arm and facial movements after focal brain lesions. *Neuropsychologia* 19: 491–503, 1981.

Kolling N, Wittmann MK, Behrens TEJ, Boorman ED, Mars RB, Rushworth MFS. Value, search, persistence and model updating in anterior cingulate cortex. *Nature Neuroscience* 19: 1280–1285, 2016. Number: 10 Publisher: Nature Publishing Group.

Komiyama T, Sato TR, O'Connor DH, Zhang YX, Huber D, Hooks BM, Gabitto M, Svoboda K. Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature* 464: 1182–1186, 2010.

Krakauer JW, Hadjiosif AM, Xu J, Wong AL, Haith AM. Motor Learning. *Comprehensive Physiology* 9: 613–663, 2019.

Kriegeskorte N, Douglas PK. Interpreting encoding and decoding models. *Current Opinion in Neurobiology* 55: 167–179, 2019.

Kropotov JD, Etlinger SC. Selection of actions in the basal ganglia–thalamocortical circuits: review and model. *International Journal of Psychophysiology* 31: 197–217, 1999.

Kurnikova A, Moore JD, Liao SM, Deschênes M, Kleinfeld D. Coordination of Orofacial Motor Actions into Exploratory Behavior by Rat. *Current Biology* 27: 688–696, 2017.

Laboy-Juárez KJ, Langberg T, Ahn S, Feldman DE. Elementary motion sequence detectors in whisker somatosensory cortex. *Nature Neuroscience* 22: 1438–1449, 2019. Number: 9 Publisher: Nature Publishing Group.

Laplane D, Talairach J, Meininger V, Bancaud J, Orgogozo JM. Clinical consequences of corticectomies involving the supplementary motor area in man. *Journal of the Neurological Sciences* 34: 301–314, 1977.

Lashley KS. The problem of serial order in behavior. In *Cerebral mechanisms in behavior; the Hixon Symposium*, pp. 112–146. Oxford, England: Wiley, 1951.

LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature* 521: 436–444, 2015. Number: 7553 Publisher: Nature Publishing Group.

Leibo JZ, d'Autume CdM, Zoran D, Amos D, Beattie C, Anderson K, Castañeda AG, Sanchez M, Green S, Gruslys A, Legg S, Hassabis D, Botvinick MM. Psychlab: A Psy-

chology Laboratory for Deep Reinforcement Learning Agents. *arXiv:180108116 [cs, q-bio]* , 2018. ArXiv: 1801.08116.

Leinweber M, Ward DR, Sobczak JM, Attinger A, Keller GB. A Sensorimotor Circuit in Mouse Cortex for Visual Flow Predictions. *Neuron* 95: 1420–1432.e5, 2017.

Li N, Chen S, Guo ZV, Chen H, Huo Y, Inagaki HK, Chen G, Davis C, Hansel D, Guo C, Svoboda K. Spatiotemporal constraints on optogenetic inactivation in cortical circuits. *eLife* 8: e48622, 2019.

Li N, Chen TW, Guo ZV, Gerfen CR, Svoboda K. A motor cortex circuit for motor planning and movement. *Nature* 519: 51–56, 2015.

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

Li Y, Lu H, Cheng Pl, Ge S, Xu H, Shi SH, Dan Y. Clonally related visual cortical neurons show similar stimulus feature selectivity. *Nature* 486: 118–121, 2012. Number: 7401 Publisher: Nature Publishing Group.

Lin LD, Murray GM, Sessle BJ. The effect of bilateral cold block of the primate face primary somatosensory cortex on the performance of trained tongue-protrusion task and biting tasks. *Journal of Neurophysiology* 70: 985–996, 1993. Publisher: American Physiological Society.

Lin LD, Murray GM, Sessle BJ. Functional properties of single neurons in the primate face primary somatosensory cortex. I. Relations with trained orofacial motor behaviors. *Journal of Neurophysiology* 71: 2377–2390, 1994. Publisher: American Physiological Society.

Liu D, Gu X, Zhu J, Zhang X, Han Z, Yan W, Cheng Q, Hao J, Fan H, Hou R, Chen Z, Chen Y, Li CT. Medial prefrontal activity during delay period contributes to learning of a working memory task. *Science* 346: 458–463, 2014.

Liu Y, Dolan RJ, Kurth-Nelson Z, Behrens TEJ. Human Replay Spontaneously Reorganizes Experience. *Cell* 178: 640–652.e14, 2019.

Logan GD. On the ability to inhibit complex movements: A stop-signal study of typewriting. *Journal of Experimental Psychology: Human Perception and Performance* 8: 778–792, 1982. Place: US Publisher: American Psychological Association.

Lu L, Cao Y, Tokita K, Heck DH, Boughter JDJ. Medial cerebellar nuclear projections and activity patterns link cerebellar output to orofacial and respiratory behavior. *Frontiers in Neural Circuits* 7, 2013. Publisher: Frontiers.

Lu X, Ashe J. Anticipatory activity in primary motor cortex codes memorized movement sequences. *Neuron* 45: 967–973, 2005.

Lu X, Matsuzawa M, Hikosaka O. A Neural Correlate of Oculomotor Sequences in Supplementary Eye Field. *Neuron* 34: 317–325, 2002.

MacDonald CJ, Meck WH, Simon SA, Nicolelis MAL. Taste-Guided Decisions Differentially Engage Neuronal Ensembles across Gustatory Cortices. *Journal of Neuroscience* 29: 11271–11282, 2009. Publisher: Society for Neuroscience Section: Articles.

Mackevicius EL, Fee MS. Building a state space for song learning. *Current Opinion in Neurobiology* 49: 59–68, 2018.

Madisen L, Mao T, Koch H, Zhuo Jm, Berenyi A, Fujisawa S, Hsu YWA, Garcia AJ, Gu X, Zanella S, Kidney J, Gu H, Mao Y, Hooks BM, Boyden ES, Buzsáki G, Ramirez JM, Jones AR, Svoboda K, Han X, Turner EE, Zeng H. A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nature Neuroscience* 15: 793–802, 2012.

Makino H, Komiyama T. Learning enhances the relative impact of top-down processing in the visual cortex. *Nature Neuroscience* 18: 1116–1122, 2015.

Martin KE, Phillips JG, Iansek R, Bradshaw JL. Inaccuracy and instability of sequential movements in Parkinson's disease. *Experimental Brain Research* 102: 131–140, 1994.

Mathis MW, Mathis A, Uchida N. Somatosensory Cortex Plays an Essential Role in Forelimb Motor Adaptation in Mice. *Neuron* 93: 1493–1503.e6, 2017.

Matsumoto N, Hanakawa T, Maki S, Graybiel AM, Kimura M. Nigrostriatal Dopamine System in Learning to Perform Sequential Motor Tasks in a Predictive Manner. *Journal of Neurophysiology* 82: 978–998, 1999. Publisher: American Physiological Society.

Mayrhofer JM, El-Boustani S, Foustoukos G, Auffret M, Tamura K, Petersen CC. Distinct Contributions of Whisker Sensory Cortex and Tongue-Jaw Motor Cortex in a Goal-Directed Sensorimotor Transformation. *Neuron* p. S0896627319306348, 2019.

McClung JR, Goldberg SJ. Organization of motoneurons in the dorsal hypoglossal nucleus that innervate the retrusor muscles of the tongue in the rat. *The Anatomical Record* 254: 222–230, 1999. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/%28SICI%291097-0185%2819990201%29254%3A2%3C222%3A%3AAID-AR8%3E3.0.CO%3B2-B.

McClung JR, Goldberg SJ. Functional anatomy of the hypoglossal innervated muscles of the rat tongue: A model for elongation and protrusion of the mammalian tongue. *The Anatomical Record* 260: 378–386, 2000. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/1097-0185%2820001201%29260%3A4%3C378%3A%3AAID-AR70%3E3.0.CO%3B2-A.

McDonald RJ, Hong NS. How does a specific learning and memory system in the mammalian brain gain control of behavior? *Hippocampus* 23: 1084–1102, 2013. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hipo.22177.

McElvain LE, Friedman B, Karten HJ, Svoboda K, Wang F, Deschênes M, Kleinfeld D.

Circuits in the Rodent Brainstem that Control Whisking in Concert with Other Orofacial Motor Actions. *Neuroscience* 368: 152–170, 2018.

Miller A. Oral and Pharyngeal Reflexes in the Mammalian Nervous System: Their Diverse Range in Complexity and the Pivotal Role of the Tongue. *Critical Reviews in Oral Biology & Medicine* 13: 409–425, 2002. Publisher: SAGE Publications Inc.

Miller KJ, Botvinick MM, Brody CD. Dorsal hippocampus contributes to model-based planning. *Nature Neuroscience* 20: 1269–1276, 2017. Number: 9 Publisher: Nature Publishing Group.

Mita A, Mushiake H, Shima K, Matsuzaka Y, Tanji J. Interval time coding by neurons in the presupplementary and supplementary motor areas. *Nature Neuroscience* 12: 502–507, 2009. Number: 4 Publisher: Nature Publishing Group.

Miyachi S, Hikosaka O, Miyashita K, Kárádi Z, Rand MK. Differential roles of monkey striatum in learning of sequential hand movement:. *Experimental Brain Research* 115: 1–5, 1997.

Mnih V, Kavukcuoglu K, Silver D, Graves A, Antonoglou I, Wierstra D, Riedmiller M. Playing Atari with Deep Reinforcement Learning. *arXiv:13125602 [cs]* , 2013. ArXiv: 1312.5602.

Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, Petersen S, Beattie C, Sadik A, Antonoglou I, King H, Kumaran D, Wierstra D, Legg S, Hassabis D. Human-level control through deep reinforcement learning. *Nature* 518: 529–533, 2015.

Moore JD, Deschênes M, Furuta T, Huber D, Smear MC, Demers M, Kleinfeld D. Hierarchy of orofacial rhythms revealed through whisking and breathing. *Nature* 497: 205–210, 2013.

Moore JD, Kleinfeld D, Wang F. How the brainstem controls orofacial behaviors comprised of rhythmic actions. *Trends in Neurosciences* 37: 370–380, 2014.

Mowery TM, Kotak VC, Sanes DH. Transient Hearing Loss Within a Critical Period Causes Persistent Changes to Cellular Properties in Adult Auditory Cortex. *Cerebral Cortex* 25: 2083–2094, 2015.

Murray GM, Lin LD, Moustafa EM, Sessle BJ. Effects of reversible inactivation by cooling of the primate face motor cortex on the performance of a trained tongue-protrusion task and a trained biting task. *Journal of Neurophysiology* 65: 511–530, 1991. Publisher: American Physiological Society.

Murray GM, Sessle BJ. Functional properties of single neurons in the face primary motor cortex of the primate. I. Input and output features of tongue motor cortex. *Journal of Neurophysiology* 67: 747–758, 1992a. Publisher: American Physiological Society.

Murray GM, Sessle BJ. Functional properties of single neurons in the face primary motor cortex of the primate. II. Relations with trained orofacial motor behavior. *Journal of Neurophysiology* 67: 759–774, 1992b. Publisher: American Physiological Society.

Murray GM, Sessle BJ. Functional properties of single neurons in the face primary motor cortex of the primate. III. Relations with different directions of trained tongue protrusion. *Journal of Neurophysiology* 67: 775–785, 1992c. Publisher: American Physiological Society.

Musall S, Urai AE, Sussillo D, Churchland AK. Harnessing behavioral diversity to understand neural computations for cognition. *Current Opinion in Neurobiology* 58: 229–238, 2019.

Mushiake H, Inase M, Tanji J. Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential

movements. *Journal of Neurophysiology* 66: 705–718, 1991. Publisher: American Physiological Society.

Navratilova Z, Godfrey KB, McNaughton BL. Grids from bands, or bands from grids? An examination of the effects of single unit contamination on grid cell firing fields. *Journal of Neurophysiology* 115: 992–1002, 2016.

Niv Y. Learning task-state representations. *Nature Neuroscience* 22: 1544–1553, 2019.

Norman DA, Shallice T. Attention to Action. In Davidson RJ, Schwartz GE, Shapiro D (eds.) *Consciousness and Self-Regulation: Advances in Research and Theory Volume 4*, pp. 1–18. Boston, MA: Springer US, 1986.

Numminen J, Schürmann M, Hiltunen J, Joensuu R, Jousmäki V, Koskinen SK, Salmelin R, Hari R. Cortical activation during a spatiotemporal tactile comparison task. *NeuroImage* 22: 815–821, 2004.

Okubo TS, Mackevicius EL, Payne HL, Lynch GF, Fee MS. Growth and splitting of neural sequences in songbird vocal development. *Nature* 528: 352–357, 2015.

Pachitariu M, Steinmetz N, Kadir S, Carandini M, D HK. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv* p. 061481, 2016. Publisher: Cold Spring Harbor Laboratory Section: New Results.

Panigrahi B, Martin K, Li Y, Graves A, Vollmer A, Olson L, Mensh B, Karpova A, Dudman J. Dopamine Is Required for the Neural Representation and Control of Movement Vigor. *Cell* 162: 1418–1430, 2015.

Parent A, Hazrati LN. Functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop. *Brain Research Reviews* 20: 91–127, 1995.

Pierrot-Deseilligny C, Rivaud S, Gaymard B, Müri R, Vermersch AI. Cortical control of saccades: Cortical Control of Saccades. *Annals of Neurology* 37: 557–567, 1995.

Preuss TM. Do Rats Have Prefrontal Cortex? The Rose-Woolsey-Akert Program Reconsidered. *Journal of Cognitive Neuroscience* 7: 1–24, 1995. Publisher: MIT Press.

Procyk E, Tanaka YL, Joseph JP. Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nature Neuroscience* 3: 502–508, 2000. Number: 5 Publisher: Nature Publishing Group.

Proske U, Gandevia SC. The proprioceptive senses: their roles in signaling body shape, body position and movement, and muscle force. *Physiological Reviews* 92: 1651–1697, 2012.

Pruszynski JA, Kurtzer I, Nashed JY, Omrani M, Brouwer B, Scott SH. Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature* 478: 387–390, 2011.

Rao RPN, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* 2: 79–87, 1999. Number: 1 Publisher: Nature Publishing Group.

Richards BA, Lillicrap TP, Beaudoin P, Bengio Y, Bogacz R, Christensen A, Clopath C, Costa RP, de Berker A, Ganguli S, Gillon CJ, Hafner D, Kepecs A, Kriegeskorte N, Latham P, Lindsay GW, Miller KD, Naud R, Pack CC, Poirazi P, Roelfsema P, Sacramento J, Saxe A, Scellier B, Schapiro AC, Senn W, Wayne G, Yamins D, Zenke F, Zylberberg J, Therien D, Kording KP. A deep learning framework for neuroscience. *Nature Neuroscience* 22: 1761–1770, 2019.

Roland PE, Larsen B, Lassen NA, Skinhoj E. Supplementary motor area and other cortical areas in organization of voluntary movements in man. *Journal of Neurophysiology* 43: 118–136, 1980. Publisher: American Physiological Society.

Rosenbaum DA. *Human motor control.* Amsterdam ; Boston, MA: Elsevier Inc, 2009, 2nd ed edn.

Rosenbaum DA, Weber RJ, Hazelett WM, Hindorff V. The parameter remapping effect in human performance: Evidence from tongue twisters and finger fumblers. *Journal of Memory and Language* 25: 710–725, 1986.

Russo AA, Bittner SR, Perkins SM, Seely JS, London BM, Lara AH, Miri A, Marshall NJ, Kohn A, Jessell TM, Abbott LF, Cunningham JP, Churchland MM. Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron* 97: 953–966.e8, 2018.

Russo AA, Khajeh R, Bittner SR, Perkins SM, Cunningham JP, Abbott LF, Churchland MM. Neural trajectories in the supplementary motor area and primary motor cortex exhibit distinct geometries, compatible with different classes of computation. preprint, Neuroscience, 2019.

Saffran JR, Kirkham NZ. Infant Statistical Learning. *Annual Review of Psychology* 69: 181–203, 2018.

Sakai K, Kitaguchi K, Hikosaka O. Chunking during human visuomotor sequence learning. *Experimental Brain Research* 152: 229–242, 2003.

Savitt JM. Bcl-x Is Required for Proper Development of the Mouse Substantia Nigra. *Journal of Neuroscience* 25: 6721–6728, 2005.

Schneider DM. Reflections of action in sensory cortex. *Current Opinion in Neurobiology* 64: 53–59, 2020.

Schneider DM, Sundararajan J, Mooney R. A cortical filter that learns to suppress the acoustic consequences of movement. *Nature* 561: 391–395, 2018. Number: 7723 Publisher: Nature Publishing Group.

Scott SH, Cluff T, Lowrey CR, Takei T. Feedback control during voluntary motor actions. *Current Opinion in Neurobiology* 33: 85–94, 2015.

Sessle BJ. Chapter 5 - Face sensorimotor cortex: Its role and neuroplasticity in the control of orofacial movements. In Gossard JP, Dubuc R, Kolta A (eds.) *Progress in Brain Research*, vol. 188 of *Breathe, Walk and Chew: The Neural Challenge: Part II*, pp. 71–82. Elsevier, 2011.

Severson KS, Xu D, Van de Loo M, Bai L, Ginty DD, O'Connor DH. Active Touch and Self-Motion Encoding by Merkel Cell-Associated Afferents. *Neuron* 94: 666–676.e9, 2017.

Shadmehr R, Smith MA, Krakauer JW. Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience* 33: 89–108, 2010.

Shea CH, Kovacs AJ, Panzer S. The Coding and Inter-Manual Transfer of Movement Sequences. *Frontiers in Psychology* 2, 2011. Publisher: Frontiers.

Shenoy KV, Sahani M, Churchland MM. Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience* 36: 337–359, 2013.

Shibasaki H, Sadato N, Lyshkow H, Yonekura Y, Honda M, Nagamine T, Suwazono S, Magata Y, Ikeda A, Miyazaki M, Fukuyama H, Asato R, Konishi J. Both primary motor cortex and supplementary motor area play an important role in complex finger movement. *Brain* 116: 1387–1398, 1993. Publisher: Oxford Academic.

Shima K, Mushiake H, Saito N, Tanji J. Role for cells in the presupplementary motor area in updating motor plans. *Proceedings of the National Academy of Sciences* 93: 8694–8698, 1996. Publisher: National Academy of Sciences Section: Research Article.

Shima K, Tanji J. Both Supplementary and Presupplementary Motor Areas Are Crucial for the Temporal Organization of Multiple Movements. *Journal of Neurophysiology* 80: 3247–3260, 1998. Publisher: American Physiological Society.

Shima K, Tanji J. Neuronal activity in the supplementary and presupplementary motor

areas for temporal organization of multiple movements. *Journal of Neurophysiology* 84: 2148–2160, 2000.

Simoncelli EP, Olshausen BA. Natural Image Statistics and Neural Representation. *Annual Review of Neuroscience* 24: 1193–1216, 2001.

Smith SL, Smith IT, Branco T, Häusser M. Dendritic spikes enhance stimulus selectivity in cortical neurons in vivo. *Nature* 503: 115–120, 2013. Number: 7474 Publisher: Nature Publishing Group.

Sofroniew NJ, Vlasov YA, Hires SA, Freeman J, Svoboda K. Neural coding in barrel cortex during whisker-guided locomotion. *eLife* 4: e12559, 2015. Publisher: eLife Sciences Publications, Ltd.

Stachenfeld KL, Botvinick MM, Gershman SJ. The hippocampus as a predictive map. *Nature Neuroscience* 20: 1643–1653, 2017.

Stanek E IV, Cheng S, Takatoh J, Han BX, Wang F. Monosynaptic premotor circuit tracing reveals neural substrates for oro-motor coordination. *eLife* 3: e02511, 2014. Publisher: eLife Sciences Publications, Ltd.

Stapleton JR. Rapid Taste Responses in the Gustatory Cortex during Licking. *Journal of Neuroscience* 26: 4126–4138, 2006.

Stavisky SD, Kao JC, Ryu SI, Shenoy KV. Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions. *Neuron* 95: 195–208.e9, 2017.

Steinmetz NA, Koch C, Harris KD, Carandini M. Challenges and opportunities for large-scale electrophysiology with Neuropixels probes. *Current Opinion in Neurobiology* 50: 92–100, 2018.

Strogatz SH. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering.* Studies in Nonlinearity. Cambridge, Mass: Perseus Books, 2001, 2. print edn. OCLC: 248691524.

Sutton RS, Barto AG. *Reinforcement Learning: An Introduction.* MIT Press, 2018. Google-Books-ID: uWV0DwAAQBAJ.

Svoboda K, Li N. Neural mechanisms of movement planning: motor cortex and beyond. *Current Opinion in Neurobiology* 49: 33–41, 2018.

Tanji J. Sequential Organization of Multiple Movements: Involvement of Cortical Motor Areas. *Annual Review of Neuroscience* 24: 631–651, 2001.

Tanji J, Shima K. Role for supplementary motor area cells in planning several movements ahead. *Nature* 371: 413–416, 1994. Number: 6496 Publisher: Nature Publishing Group.

Telgen S, Parvin D, Diedrichsen J. Mirror Reversal and Visual Rotation Are Learned and Consolidated via Separate Mechanisms: Recalibrating or Learning De Novo. *The Journal of Neuroscience* 34: 13768–13779, 2014.

Thompson PD, Berardelli A, Rothwell JC, Day BL, Dick JPR, Benecke R, Marsden CD. THE COEXISTENCE OF BRADYKINESIA AND CHOREA IN HUNTINGTON'S DISEASE AND ITS IMPLICATIONS FOR THEORIES OF BASAL GANGLIA CONTROL OF MOVEMENT. *Brain* 111: 223–244, 1988. Publisher: Oxford Academic.

Travers JB, Dinardo LA, Karimnamazi H. Motor and Premotor Mechanisms of Licking. *Neuroscience & Biobehavioral Reviews* 21: 631–647, 1997.

Travers JB, Herman K, Travers SP. Suppression of third ventricular NPY-elicited feeding following medullary reticular formation infusions of muscimol. *Behavioral Neuroscience* 124: 225–233, 2010. Place: US Publisher: American Psychological Association.

Treuting PM, Dintzis SM, Montine KS. *Comparative Anatomy and Histology: A Mouse, Rat, and Human Atlas.* Academic Press, 2017. Google-Books-ID: FgBQCwAAQBAJ.

Van Den Bercken JHL, Cools AR. Evidence for a role of the caudate nucleus in the sequential organization of behaviour. *Behavioural Brain Research* 4: 319–337, 1982.

Vong L, Ye C, Yang Z, Choi B, Chua S, Lowell BB. Leptin action on GABAergic neurons prevents obesity and reduces inhibitory tone to POMC neurons. *Neuron* 71: 142–154, 2011.

Watabe-Uchida M, Eshel N, Uchida N. Neural Circuitry of Reward Prediction Error. *Annual Review of Neuroscience* 40: 373–394, 2017. _eprint: https://doi.org/10.1146/annurev-neuro-072116-031109.

Welzl H, Bureš J. Lick-synchronized breathing in rats. *Physiology & Behavior* 18: 751–753, 1977.

Williams AH, Kim TH, Wang F, Vyas S, Ryu SI, Shenoy KV, Schnitzer M, Kolda TG, Ganguli S. Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis. *Neuron* 98: 1099–1115.e8, 2018.

Wilson R, Takahashi Y, Schoenbaum G, Niv Y. Orbitofrontal Cortex as a Cognitive Map of Task Space. *Neuron* 81: 267–279, 2014.

Xu D, Chen Y, Delgado AM, Hughes NC, Dong M, Zhang L, O'Connor DH. A functional cortical network for sensorimotor sequence generation. *bioRxiv* p. 783050, 2020. Publisher: Cold Spring Harbor Laboratory Section: New Results.

Xu S, Jiang W, Poo Mm, Dan Y. Activity recall in a visual cortical ensemble. *Nature Neuroscience* 15: 449–455, 2012.

Yamashita Y, Tani J. Emergence of Functional Hierarchy in a Multiple Timescale Neural Network Model: A Humanoid Robot Experiment. *PLoS Computational Biology* 4: e1000220, 2008.

Yang GR, Joglekar MR, Song HF, Newsome WT, Wang XJ. Task representations in neural networks trained to perform many cognitive tasks. *Nature Neuroscience* 22: 297–306, 2019. Number: 2 Publisher: Nature Publishing Group.

Yoshida T, Katz DB. Control of Prestimulus Activity Related to Improved Sensory Coding within a Discrimination Task. *Journal of Neuroscience* 31: 4101–4112, 2011. Publisher: Society for Neuroscience Section: Articles.

Zhang YS, Ghazanfar AA. A Hierarchy of Autonomous Systems for Vocal Production. *Trends in Neurosciences* 43: 115–126, 2020.

Zhao S, Ting JT, Atallah HE, Qiu L, Tan J, Gloss B, Augustine GJ, Deisseroth K, Luo M, Graybiel AM, Feng G. Cell type–specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. *Nature Methods* 8: 745–752, 2011.

Zhou X, Wang L, Hasegawa H, Amin P, Han BX, Kaneko S, He Y, Wang F. Deletion of PIK3C3/Vps34 in sensory neurons causes rapid neurodegeneration by disrupting the endosomal but not the autophagic pathway. *Proceedings of the National Academy of Sciences* 107: 9424–9429, 2010.

Zmarz P, Keller GB. Mismatch Receptive Fields in Mouse Visual Cortex. *Neuron* 92: 766–772, 2016.

Curriculum Vitae

Duo Xu

EDUCATION

Aug 2014 – present   **Johns Hopkins University School of Medicine**

Program: Neuroscience Graduate Training Program, Solomon H. Snyder Department of Neuroscience

Status: Ph.D. candidate (expected award date: May 15, 2019)

Sep 2008 – Jun 2012   **China Agricultural University**

Major: Honors Program Biological Sciences

Degree: Bachelor of Science

RESEARCH EXPERIENCE

Aug 2014 – present   **Department of Neuroscience, Johns Hopkins University**

Position: Ph.D. student

Supervisor: Daniel H. O'Connor, Associate Professor

- Developing a novel sensorimotor sequence task for mice and characterizing a functional cortical network for sequence generation and control
- Investigation of the touch and self-motion coding by Merkel cell-associated primary afferents

Jun 2011 – Jul 2014   **Institute of Neuroscience, Chinese Academy of Sciences**

Position: Research assistant

Supervisor: Mu-ming Poo, Senior Investigator and Director

- Imaging BDNF secretion from adult-born hippocampal neurons
- Developing an automated control system for a reaching behavior in mice
- Designing an intracranial chamber and an artificial dura for imaging and electrophysiology in macaque

College of Biological Sciences, China Agricultural University

Position: Undergraduate student

Nov 2010 – Jun 2011   Supervisor: Sheng Cui, Professor

- Investigation of the potential role of miRNA-218 in the developing spinal cord

May – Nov 2010   Supervisor: Ning Li, Professor

- Analyzing the efficiency of piggyBac transposon transgenic method in chick

May – Nov 2009   Supervisor: Baoqing Wang, Associate Professor

- Characterization of the stimuli which induce immediate escape behavior of crabs

PUBLICATIONS & INNOVATION

- **Xu D**, Chen Y, Delgado AM, Hughes NC, Zhang L, Dong M, O'Connor DH. A functional cortical network for sensorimotor sequence generation. *bioRxiv* 783050, 2019.
- Severson KS, **Xu D**, Yang H, O'Connor DH. Coding of whisker motion across the mouse face. *eLife*, 2019.

- Severson KS*, **Xu D***, Loo MV, Bai L, Ginty DD, O'Connor DH. Active Touch and Self-Motion Encoding by Merkel Cell-Associated Afferents. *Neuron*. (* equal contribution), 2017.
- Shang C, **Xu D**, Poo M. Chinese translation of the book, *The Future of the Brain* (Steve Rose), 2016.
- Wang L, Chang X, She L, **Xu D**, Huang W, Poo M. Autocrine Action of BDNF on Dendrite Development of Adult-Born Hippocampal Neurons. *Journal of Neuroscience*, 2015.
- ManySyncX, an open source file synchronization software, 2013.

PRESENTATIONS & POSTERS

- **Xu D**, Chen Y, Delgado AM, O'Connor DH. Sensorimotor control of complex tongue movement sequences in mouse. Nanosymposium, Society for Neuroscience Annual Meeting, Chicago, IL, 2019.
- A functional cortical network for sensorimotor sequence generation. Kavli NDI breakfast talk, 2019.
- **Xu D**, Chen Y, Delgado AM, O'Connor DH. Sensory-motor sequence generation and learning of tongue movements. Poster, Society for Neuroscience Annual Meeting, San Diego, CA, 2018.

HONORS AND AWARDS

- Poster Award, Johns Hopkins Neuroscience Department Annual Retreat   2019
- Excellence scholarship   2008 – 2011

SKILLS

Data analysis (MATLAB and Python)	Behavioral task development
Software development (C# and MATLAB)	Extracellular recording
Embedded system development (C++)	Two-photon imaging
Mechanical design (Solidworks)	Whole-cell recording (*in-vitro*)
Electrical circuit design (Eagle)	Cell and tissue culture
Deep learning (TensorFlow and MATLAB)	*In-ovo* electroporation